POSET STRUCTURES ON $(m+2)$-ANGULATIONS AND POLYNOMIAL BASES OF THE QUOTIENT BY G^m-QUASISYMMETRIC FUNCTIONS

FRÉDÉRIC CHAPOTON AND JEAN-CHRISTOPHE AVAL

ABSTRACT. For integers $m, n \geq 1$, we describe a bijection sending dissections of the $(mn+2)$-regular polygon into $(m+2)$-sided polygons to a new basis of the quotient of the polynomial algebra in mn variables by an ideal generated by some kind of higher quasi-symmetric functions. We show that divisibility of the basis elements corresponds to a new partial order on dissections, which is studied in some detail.

1. INTRODUCTION

Let $m \geq 1$ be an integer. For every integer $n \geq 1$, we define a simple poset structure $\mathcal{P}_{m,n}$ on the set of $(m+2)$-angulations of a $(mn+2)$-gon. This generalizes the construction by Pallo on triangulations [Pal03], which is closely related to the Tamari lattice.

Quasisymmetric functions, a generalisation of symmetric functions, were introduced in [Ges84] and are now classical in algebraic combinatorics. Some higher analogues were introduced in [Poi98] and further studied in [NT10, BH08]. We recall some results about quotients of polynomials rings by higher quasi-symmetric functions, first obtained for quasisymmetric functions in [AB03, ABB04] and extended to the higher case in [Ava07].

We then show that the poset $\mathcal{P}_{m,n}$ is isomorphic to the divisibility poset of a new particular basis of the quotient of the polynomial ring in m sets of n variables by the ideal generated by G^m-quasisymmetric functions without constant term. Our description of a new basis builds upon the basis indexed by m-Dyck paths that was introduced for general m in [Ava07].

For $m = 1$, the posets $\mathcal{P}_{1,n}$ were introduced by Pallo in [Pal03] and further studied in [CSS12, CSS14], and the basis indexed by triangulations was defined in [Cha05], but the connection between them is new.

The last two sections of the article are devoted to some enumerative results on the posets (enumeration of intervals, rank generating function) and to a recursive description of the intervals as distributive lattices of orders ideals of forests.

2. A POSET STRUCTURE ON M-ANGULATIONS

Let $m \geq 1$ be an integer. For the sake of readability, we will use the expression M-angulation instead of $(m+2)$-angulation. An M-angulation of a regular convex polygon is a set of diagonals that cut the polygon into $(m+2)$-sided polygonal regions.

We consider the set $\mathcal{Q}_{m,n}$ of M-angulations of a $(mn+2)$-gon. Every element of $\mathcal{Q}_{m,n}$ contains n regions, separated by $n-1$ diagonal edges. The cardinality of $\mathcal{Q}_{m,n}$ is given by $\frac{1}{mn+1}\binom{(m+1)n}{n}$, the number of $(m+1)$-ary planar rooted trees with n inner vertices (often called a Fuss-Catalan number). A simple bijection between these two classes of

objects is given by planar duality. Some elements of $\mathcal{Q}_{2,7}$ are shown in Figures 1 and 2.

FIGURE 1. A quadrangulation of the 16-gon

To define a poset structure on $\mathcal{Q}_{m,n}$, we fix a particular element Q_0, which is a fan (every diagonal edge involves a fixed vertex, denoted 0 and called the *apex*, see Figure 2).


FIGURE 2. The fan Q_0 for $m = 2$ and $n = 7$

We consider the following order relation. An element Q of $\mathcal{Q}_{m,n}$ is covered in $\mathcal{P}_{m,n}$ by the M-angulations obtained by flipping one of its diagonal edges included in Q_0. Here flipping means removing this diagonal edge and replacing it by another diagonal edge cutting again the $(2m+2)$-gon created by the removal into two $(m+2)$-gons. Note that any diagonal edge may be flipped in exactly m different ways. As an example the poset $\mathcal{P}_{2,3}$ is shown on Figure 3. Note also that once a diagonal x has been flipped in a cover relation $x \leq y$, the chosen new diagonal will belong to all z greater than y, because it does not belong to Q_0.

This poset is clearly graded, the rank of an element Q being $n-1$ minus the number of diagonal edges shared by Q_0 and Q. This is also the number of diagonal edges in Q that are not in Q_0.

It follows from the description of the coverings that an element of rank r is covered by exactly $m(n-1-r)$ elements. This implies the formula $m^{n-1}(n-1)!$ for the number of maximal chains in $\mathcal{P}_{m,n}$.

This poset has Q_0 as unique minimum. The fact that Q_0 is smaller than all M-angulations follows from the next lemma by induction on the number of diagonals that

FIGURE 3. The poset $\mathcal{P}_{2,3}$

are both in Q and Q_0. This lemma implies that unless Q is Q_0, there is a cover relation $Q' \lhd Q$ where Q' has one more diagonal in common with Q_0.

Lemma 1. *Let Q be distinct from Q_0. Then there is always at least one diagonal d_0 in Q_0 that cuts exactly one diagonal of Q.*

Proof. One can assume without loss of generality that Q and Q_0 have no common diagonal, otherwise one can find d_0 (by induction) inside one of the parts cut by the common diagonals. In at least one of these parts the restriction of Q must differ from the restriction of Q_0, because Q is not Q_0.

Let us label all the vertices of the regular polygon counter-clockwise by integers, starting from the vertex 0 which is the apex of Q_0.

Because Q and Q_0 have no common diagonal, there exists a unique region R_0 of Q that contains the vertex 0 in its boundary. Removing R_0 from the ambient polygon, we get one or more convex polygons, all of them with $km + 2$ vertices for some k. Let us choose one of these polygons, and let Q' be its M-angulation obtained from Q by restriction. Let R' be the unique region of Q' which is adjacent to R_0.

Excluding minimal and maximal indices, vertices of Q' form a sequence of mk vertices numbered consecutively, in which every residue class modulo m is represented exactly k times.

Because of the tree-like structure of M-angulations, one can build the M-angulation Q' by successive additions of M-angles, starting from the region R'. Going backwards, one can go from Q' to R' by removal of M-angles on the boundary.

Using this leaf-removal induction, one can then prove that the boundary of R' (excluding minimal and maximal indices) contains exactly one representative of every residue class modulo m.

Note now that, in Q_0, the vertices linked to the apex 0 by a diagonal form a residue class modulo m, when numbered in the same way as the vertices of Q.

It follows that exactly one of the vertices of R' (excluding minimal and maximal indices) is the end of a diagonal d_0 in Q_0. The diagonal d_0 crosses only one diagonal of Q, namely the diagonal separating the regions R_0 and R'. $\qquad\square$

The maximal elements in the poset $\mathcal{P}_{m,n}$ are the M-angulations that have no common diagonal with Q_0. We will call them *final M-angulations*.

Remark. There are several interesting existing families of posets on objects in bijection with $(m+2)$-angulations, including m-Tamari lattices [BPR12] and m-Cambrian lattices of type A [STW15]. The posets introduced here seem to be new.

3. m-ANALOGUES OF B-QUASISYMMETRIC FUNCTIONS

In [BH08] and [NT10], the ring of B-quasisymmetric functions was introduced as the graded dual Hopf algebra of the analogue in type B of Solomon's descent algebra. This ring is contained in the polynomial ring in two sets of n variables.

In [Ava07], a quotient ring of this polynomial ring by an ideal of B-quasisymmetric functions was studied. Moreover, for every integer $m \geq 1$, an analogue of the ring of B-quasisymmetric functions and an analogue of the quotient ring were also defined and studied, involving m sets of n variables.

We refer to [Ava07] for the original motivations of the study of these rings and for the proof of the results that we will use. Let us now summarize the results of [Ava07] in their most general form.

Let us denote by $\mathbb{M} = \{x, y, z, \ldots, \omega\}$ a set of m distinct letters, endowed with a total order $x < y < z < \cdots < \omega$. We will mostly illustrate our constructions with the cases $m = 2$ and $m = 3$, therefore using only the letters x, y, z.

Let us start with polynomials in the union of m pairwise disjoint alphabets of n variables each: $X_n = x_1, \ldots, x_n$, $Y_n = y_1, \ldots, y_n$, up to $\Omega_n = \omega_1, \ldots, \omega_n$. Denote this polynomial ring by $\mathbb{Q}[X_n, Y_n, \ldots, \Omega_n]$. Inside this polynomial ring, one can define a space of G^m-quasisymmetric functions, which reduces to the classical quasisymmetric functions for $m = 1$.

In [Ava07], a Gröbner basis for the ideal $\mathcal{J}_{m,n}$ generated by constant-term-free G^m-quasisymmetric functions was described, and from that was deduced a monomial basis for the quotient $\mathcal{R}_{m,n}$ of the polynomial ring $\mathbb{Q}[X_n, Y_n, \ldots, \Omega_n]$ by $\mathcal{J}_{m,n}$.

This monomial basis for the quotient $\mathcal{R}_{m,n}$ is indexed by m-Dyck paths, which gives the dimension formula $\dim \mathcal{R}_{m,n} = \frac{1}{mn+1}\binom{(m+1)n}{n}$.

3.1. Definitions.
For the definitions below, we follow [BH08], with some minor differences, for the sake of simplicity of the computations we will have to make. The main change is to describe directly the general case for any $m \geq 1$ and not the special case $m = 2$.

An *m-vector* of size n is a vector $v = (v_1, v_2, \ldots, v_{mn-1}, v_{mn})$ of length mn with entries in $\mathbb{N}$. One must think of m-vectors as the concatenation of n sequences of length m. An *m-composition* is an m-vector in which there is no sequence of m consecutive zeros.

The integer n is called the *size* of v. The *weight* of v is by definition the m-tuple $(w_1, \ldots, w_m)$ where $w_j = \sum_{i=0}^{n-1} v_{mi+j}$. We also set $|v| = \sum_{i=1}^{mn} v_i$. For example, $(1, 0, 2, 1, 0, 2, 3, 0)$ is a 2-composition of size 4, and of weight $(6, 3)$.

To make notations lighter, we shall sometimes write m-vectors or m-compositions as words with bars, where the bars separate the word into n blocks of length m. For example, $10|21|02|30$ stands for the 2-vector $(1, 0, 2, 1, 0, 2, 3, 0)$ (see also the following definition).

Let us now define the *fundamental G^m-quasisymmetric polynomials*, indexed by m-compositions.

Let $c = (c_1, \dots, c_{mn})$ be an m-composition. One can decompose c as a concatenation of n blocks of m integers. To c, we associate a word w_c in the alphabet $\{x, y, z, \dots, \omega\}$, defined as the concatenation of the words $x^{b_1} y^{b_2} \dots \omega^{b_m}$ over all blocks $b = (b_1, \dots, b_m)$ of c. For example, for the 3-composition $010|201$, we obtain $w_c = yxxz$ (powers are written as repeated letters).

Then the fundamental G^m-quasisymmetric polynomial of index c is

$$F_c = \sum_i \prod_{t \in w_c} t_{i(t)},$$

where the sum is taken over all maps i from the sequence of letters of the word w_c to the set $\{1, \dots, n\}$ such that i is weakly increasing inside every block of c and strictly increasing between two blocks.

Let us give some examples for $m = 2$:

$$F_{12} = \sum_{i \leq j \leq k} x_i y_j y_k, \tag{3.1}$$

$$F_{02|10} = \sum_{i \leq j < k} y_i y_j x_k. \tag{3.2}$$

It is clear from the definition that the multidegree (i.e., the m-tuple (degree in x, degree in y, ..., degree in ω)) of F_c in $\mathbb{Q}[X_n, Y_n, \dots, \Omega_n]$ is the weight of c. If the size of c is greater than n, we set $F_c = 0$.

The space of G^m-quasisymmetric polynomials, denoted by $QSym_n(G^m)$ is the vector subspace of the ring $\mathbb{Q}[X_n, Y_n, \dots, \Omega_n]$ generated by the F_c, for all m-compositions c.

We denote the ideal $\langle QSym_n(G^m)^+ \rangle$ generated by G^m-quasisymmetric polynomials with zero constant term by $\mathcal{J}_{m,n}$.

3.2. A monomial basis for $\mathcal{R}_{m,n}$.
Let $v = (v_1, v_2, \dots, v_{mn-1}, v_{mn})$ be an m-vector of size n. We associate to v a path $\pi(v)$ in the plane $\mathbb{N} \times \mathbb{N}$, with steps $(0, 1)$ (up step) or $(m, 0)$ (right step). We start from $(0, 0)$ and for each entry v_i (read from left to right) add v_i right steps $(m, 0)$ followed by one up step $(0, 1)$. This clearly defines a bijection between m-vectors of size n and such paths of height mn.

As an example, the path associated to the 2-vector $(1, 0, 1, 2, 0, 0, 1, 1)$ is

If the path $\pi(v)$ associated with an m-vector v always remains above the diagonal $x = y$, we call this path an *m-Dyck path*, and say that the corresponding m-vector v is an *m-Dyck vector*.

Being an m-Dyck vector is equivalent to the condition that, for any $1 \le \ell \le mn$, one has

$$m(v_1 + v_2 + \cdots + v_\ell) < \ell.$$

For v an m-vector (of length mn), we denote by $\mathcal{A}_v$ the monomial

$$\mathcal{A}_v = (x_1^{v_1} y_1^{v_2} \cdots \omega_1^{v_m})(x_2^{v_{m+1}} y_2^{v_{m+2}} \cdots \omega_2^{v_{2m}}) \cdots (x_n^{v_{m(n-1)+1}} \cdots \omega_n^{v_{mn}}).$$

This clearly defines a bijection between m-vectors and all monomials in the polynomial ring $\mathbb{Q}[X_n, Y_n, \ldots, \Omega_n]$.

For example, the monomial associated to the 2-vector $(1, 0, 1, 2, 0, 0, 1, 1)$ is $x_1 x_2 y_2^2 x_4 y_4$. The following result was proved in [Ava07, Th. 5.1].

Proposition 2. *The set $\mathcal{B}_{m,n}$ of monomials $\mathcal{A}_v$ for v varying over m-Dyck vectors of size n is a basis for the space $\mathcal{R}_{m,n} = \mathbb{Q}[X_n, Y_n, \ldots, \Omega_n]/\mathcal{J}_{m,n}$.*

4. A bijection between M-angulations and m-Dyck paths

We assign to each vertex of the $(mn + 2)$-gon (except the vertex 0) a letter in $\mathbb{M}$ as follows. The vertices are labelled by repeating the sequence $x, y, z, \ldots, \omega$ in counterclockwise order around the polygon in such a way that the final vertex just before the vertex 0 receives the last letter ω of $\mathbb{M}$. It follows that the first vertex just after the vertex 0 also receives the letter ω.

See Figure 4 for an illustration of this labelling when $m = 2$ with the ordered set of letters $x < y$.


FIGURE 4. The labelling of vertices by x and y letters ($m = 2$)

We also label the inner diagonals of Q_0 from 1 to $n - 1$ in counter-clockwise order.

Then we consider an M-angulation Q. To any diagonal d of Q, we associate a polynomial m_d. If d coincides with a diagonal of Q_0, we set $m_d = 1$. Otherwise, and reading counter-clockwise, d starts from a vertex labelled by the letter u, intersects consecutive diagonals of Q_0 labelled from i to j, and ends at a vertex labelled by the letter v. Then we set $m_d = v_{j+1} - u_i$. We then associate to Q the polynomial P_Q defined as the product of m_d over its diagonals.

As an example, Figure 5 shows the polynomials associated to the quadrangulations of Figure 3, in the corresponding positions.

$$\begin{matrix} \begin{pmatrix} x_3 - y_2 \\ y_3 - x_1 \end{pmatrix} & \begin{pmatrix} x_3 - y_2 \\ x_3 - y_1 \end{pmatrix} & \begin{pmatrix} y_3 - x_2 \\ y_3 - x_1 \end{pmatrix} & \begin{pmatrix} y_3 - x_2 \\ x_2 - y_1 \end{pmatrix} & \begin{pmatrix} x_2 - y_1 \\ x_3 - y_1 \end{pmatrix} & \begin{pmatrix} y_2 - x_1 \\ x_3 - y_1 \end{pmatrix} & \begin{pmatrix} y_2 - x_1 \\ y_3 - x_1 \end{pmatrix} \end{matrix}$$

$$(x_3 - y_2) \qquad (y_3 - x_2) \qquad (x_2 - y_1) \qquad (y_2 - x_1)$$

$$1$$

FIGURE 5. The polynomials associated to the quadrangulations of Figure 3

To deal with leading terms of polynomials, we will use the lexicographic order induced by the ordering of the variables:

$$x_1 < y_1 < \cdots < \omega_1 < x_2 < y_2 < \cdots < x_n < y_n < \cdots < \omega_n. \qquad (4.1)$$

The lexicographic order is defined on monomials as follows: $\mathcal{A}_v <_{\text{lex}} \mathcal{A}_w$ if and only if the last non-zero entry of $v - w$ (componentwise) is negative.

Note that the leading monomial of the polynomial P_Q attached to an M-angulation is easily described: in every binomial factor of the form $v_{j+1} - u_i$ with $j + 1 > i$, keep only the monomial v_{j+1}.

Proposition 3. *The set of leading monomials of the polynomials P_Q when Q varies over the set of M-angulations coincides with the monomial basis $\mathcal{B}_{m,n}$.*

Proof. We need a bijection Φ between M-angulations Q and m-Dyck paths (or rather m-Dyck vectors v) such that the leading monomial of P_Q is equal to $\mathcal{A}_{\Phi(Q)}$.

The idea to define Φ is to compose the map which sends an m-angulation Q to the leading monomial of P_Q with the bijection between monomials and m-Dyck paths described in Section 3.2.

Let us instead define the reverse bijection Ψ.

Let v be an m-Dyck vector of size n. We start from the empty set of diagonals on the $(mn + 2)$-gon. We shall iteratively add diagonals. Let D denote the current set of diagonals, under construction. We read the m-vector v from left to right. To any non-zero entry c associated to variable z_k, we add a fan with c diagonals to D as follows. Let t be the letter before z in the cyclic order $x < y < \cdots < \omega < x$. The (common) ending point of the added fan is the first vertex labelled z that comes counter-clockwise after the diagonal $k - 1$ of Q_0, and the starting points are the c last vertices labelled t that are available before the ending point. Here, being available means to be not separated from the ending point by a diagonal already in D.

See Figure 6 for an example of this construction.

FIGURE 6. The quadrangulation Q associated to the monomial $x_5\, y_5^3\, y_7$
with $P_Q = (x_5 - y_4)(y_5 - x_3)(y_5 - x_2)(y_5 - x_1)(y_7 - x_6)$.

Now the key point is that, when adding such a fan, the starting points are always strictly between the apex 0 and the ending point in counter-clockwise order. This is because of the m-Dyck word property of v: for any $1 \leq \ell \leq mn$, we have

$$m(v_1 + v_2 + \cdots + v_\ell) < \ell.$$

Indeed, when numbering the vertices of the polygon counter-clockwise from the apex 0 to $mn + 1$, the index of the first vertex t in counter-clockwise order for the fan added at step ℓ is exactly $(1 + \ell) - m(v_1 + v_2 + \cdots + v_\ell) - 1$.

The leading term of the product of binomials m_d attached to the diagonals in the fan just added is therefore the monomial z_k^c, by definition of the lexicographic order.

At the end of the process, there remains only to add some diagonals of Q_0 to obtain an M-angulation. There is a unique way to do that. This does not change the product of m_d over all diagonals.

We have therefore defined a map Ψ from m-Dyck words to M-angulations. It has the property that the leading monomial of $P_{\Psi(v)}$ is equal to $\mathcal{A}_v$. Since we can recover the m-Dyck word v from $\mathcal{A}_v$, this map is injective.

This proves that Ψ is an injection between sets of equal cardinality, hence a bijection.
□

The reverse bijection Φ can be described as first removing all diagonals in Q_0 from Q, then proceeding by successive removals of fans.

Proposition 4. *The set of polynomials P_Q (when Q varies over the set of M-angulations) is a basis of the space $\mathcal{R}_{m,n}$. Endowed with the poset structure given by divisibility, this set is isomorphic to $\mathcal{P}_{m,n}$.*

Proof. By the results of [Ava07], any set of polynomials whose leading monomials coincide with $\mathcal{B}_{m,n}$ (which is the set of monomials that do not belong to the set of leading monomials of the ideal $J_{m,n}$) is a basis of the quotient.

The map $Q \mapsto P_Q$ is a bijection, with inverse obtained from the factorisation of the polynomial. The flip in Q of a diagonal of Q_0 corresponds to the multiplication of P_Q by a linear binomial. $\square$

This new basis has therefore the interesting property that its elements display in an explicit way part of the multiplication table of the quotient, and are moreover related by a very simple two-way bijection with combinatorial objects. Moreover, it has an additional symmetry property, as explained below.

Remark. the fan Q_0 has an obvious involutive symmetry that fixes the apex 0 and flips the ambient polygon. Given the labelling of the polygon by $\mathbb{M} = \{x, y, \ldots, \omega\}$, this involution acts on variables by simultaneous substitution of letters

$$\omega, x, y, z, \ldots \longleftrightarrow \omega, \ldots, z, y, x,$$

and renumbering of indices $i \leftrightarrow n + 1 - i$. By construction, the set of polynomials P_Q is sent to itself (up to signs) by this involution acting on variables.

5. ENUMERATIVE ASPECTS

5.1. **Recursive description.** Let $T_m = \sum_{n \geq 1} \#\mathcal{Q}_{m,n} x^n$ be the generating function for M-angulations according to their size n. Note that the power of x is chosen to correspond to the number of regions.

There is a classical recursive decomposition of M-angulations, which describes them as an M-angle with an M-angulation (or nothing) grafted on all sides but one. This implies that

$$T_m = x(1 + T_m)^{m+1}. \tag{5.1}$$

Recall that the final M-angulations are those which do not contain any diagonal of Q_0. They are the maximal elements of the posets $\mathcal{P}_{m,n}$.

Let F_m be the generating function for final M-angulations according to the number of regions. These objects can be decomposed as one M-angle (the unique region having 0 in its boundary), on which one can graft any M-angulation on all sides that do not contain the vertex 0. We obtain

$$F_m = x(1 + T_m)^m = T_m/(1 + T_m). \tag{5.2}$$

Let us now describe another simple decomposition of M-angulations, corresponding to the rightmost expression in the equation (5.2).

To every M-angulation Q, one can associate a list $\mathcal{L}(Q)$ of final M-angulations, obtained by cutting Q along its initial diagonals. Let us assume that these pieces are listed in counter-clockwise order.

Proposition 5. *Sending Q to $\mathcal{L}(Q)$ defines a bijection from M-angulations to lists of final M-angulations.*

The inverse bijection is very simple: given a list of final M-angulations, we glue them back along their sides into one single M-angulation.

This inverse bijection from $\mathcal{L}(Q)$ to Q can also be interpreted in the following way, that will be useful later. Given a list of k final M-angulations, we consider the M-angulation Q_0 with k regions. We then replace the regions of Q_0 by the final M-angulations, in counter-clockwise order. In the resulting M-angulation Q, the initial Q_0 can be identified with the union of all regions that are adjacent with the vertex 0.


FIGURE 7. Illustration of construction $\mathcal{L}$

5.2. **Rank generating function.** Let G_m be the generating function for the elements of all posets $\mathcal{P}_{m,n}$, according to their size n and their rank, namely

$$G_m = \sum_{n \geq 1} x^n \sum_{Q \in \mathcal{P}_{m,n}} z^{\mathrm{rk}\, Q}. \tag{5.3}$$

The generating function G_m satisfies

$$G_m = \frac{F_m(zx)/z}{1 - F_m(zx)/z} \tag{5.4}$$

where F_m is the generating function for the final elements. Indeed, any M-angulation can be written uniquely as a list of final M-angulations by cutting along the diagonal edges shared with Q_0. The rank parameter is multiplicative along this decomposition. Moreover, the rank of final M-angulations is just their size minus 1, so that the generating function of final M-angulations according to size and rank is just $F_m(zx)/z$.

Using Lagrange inversion followed by a simple summation of binomial coefficients, one can deduce from (5.1), (5.2), and (5.4) that the rank generating function of $\mathcal{P}_{m,n}$ is given by

$$\sum_{k=0}^{n-1} \frac{n-k}{n} \binom{mn+k-1}{k} z^k. \tag{5.5}$$

For $m = 1$, this enumeration according to rank was already done in [CSS14, § 3].

5.3. **Decomposition of intervals.** Let us now study the intervals in the posets $\mathcal{P}_{m,n}$.

An interval A is a pair of elements (A^-, A^+) in $\mathcal{P}_{m,n}$ that satisfies $A^- \leq A^+$. In every finite poset, the number of intervals is also the dimension of the incidence algebra.

Let us call an interval A *initial* if its lower bound A^- is Q_0. The set of initial intervals can be identified with the set $\mathcal{Q}_{m,n}$ of M-angulations in $\mathcal{P}_{m,n}$. Indeed, Q_0 is smaller than all M-angulations by Lemma 1.

Let $\mathcal{I}_{m,n}$ be the set of intervals in $\mathcal{P}_{m,n}$. Let $I_m = \sum_{n \geq 1} \#\mathcal{I}_{m,n} x^n$ be the generating function for intervals.

We can get a recursive decomposition of intervals, similar to the previous decomposition of M-angulations. To this end, we need the following construction.

Suppose that we have an M-angulation B_0 with k regions, and a list of k final M-angulations $B_1, \ldots, B_k$. From these data, we can build an M-angulation $G(B_0; B_1, \ldots, B_k)$ as follows. First build the M-angulation $\mathcal{L}^{-1}(B_1, \ldots, B_k)$. Removal of its $k-1$ initial diagonals creates a region with $mk+2$ sides. Place B_0 inside this region to define the M-angulation $G(B_0; B_1, \ldots, B_k)$. See Figure 8 for an example.

If B_0 is Q_0, the construction G is just the inverse of the $\mathcal{L}$ map.

FIGURE 8. B_0, B_1, B_2 and $G(B_0; B_1, B_2)$

Given an interval A, we can apply the map $\mathcal{L}$ to its bottom element A^-. This gives a list of final M-angulations $A_1^-, \ldots, A_k^-$, such that

$$A^- = G(Q_0; A_1^-, \ldots, A_k^-). \qquad (5.6)$$

Proposition 6. *There exists a unique M-angulation A_0^+ such that*

$$A^+ = G(A_0^+; A_1^-, \ldots, A_k^-). \qquad (5.7)$$

Proof. Uniqueness is clear by the definition of the construction G.

Existence is proved by induction on the difference of initial diagonals in A^- and A^+. If $A^+ = A^-$, then the only possible choice is $A_0^+ = Q_0$.

Otherwise, let us consider a cover relation $A^- \leq Q' \lessdot Q''$. Assume that $Q' = G(Q_0'; A_1^-, \ldots, A_k^-)$ by induction. Since the cover relation flips an initial diagonal d and does not change the other ones, the diagonal d must in fact belong to Q_0'. Therefore, we can flip it in Q_0' to get Q_0'' such that $Q'' = G(Q_0''; A_1^-, \ldots, A_k^-)$. $\qquad\square$

Keeping the same notations, we also have the following result.

Proposition 7. *Every element of the interval $[A^-, A^+]$ can be uniquely written as $G(Q; A_1^-, \ldots, A_k^-)$ for some Q in $[Q_0, A_0^+]$. The interval A is isomorphic to the interval $[Q_0, A_0^+]$.*

Proof. This result follows from the proof of Proposition 6. In fact, more is true: all elements greater than A^- can be uniquely written in the form $G(Q; A_1^-, \ldots, A_k^-)$ for some Q, and this bijection identifies the upper ideal of A^- with a smaller poset of type $\mathcal{P}$. $\qquad\square$

Proposition 6 implies the following decomposition.

Proposition 8. *The map sending an interval A to the pair $(A_0^+, \mathcal{L}(A^-))$ defines a bijection between intervals and pairs $(B_0, (B_1, \ldots, B_k))$, where B_0 is an M-angulation with k regions, and $B_1, \ldots, B_k$ are final M-angulations.*

The inverse bijection is given by

$$A^- = G(Q_0; B_1, \ldots, B_k) \quad \text{and} \quad A^+ = G(B_0; B_1, \ldots, B_k). \qquad (5.8)$$

Corollary 9. *The generating function for intervals can be expressed in terms of the generating functions for M-angulations and for final M-angulations as*

$$I_m = T_m(F_m) = T_m(x(1 + T_m)^m). \tag{5.9}$$

The second equality follows from (5.2).

6. Isomorphism types of intervals

The aim of this section is to give a description of the intervals as posets, and to prove in particular that they are distributive lattices, whence the values of the Möbius function lie in $\{-1, 0, 1\}$.

By Proposition 7, every interval is isomorphic to an initial interval. It is therefore enough to study initial intervals.

Let us call an interval *initial-final* if its minimum is Q_0 and its maximum is a final quadrangulation.

Proposition 10. *Every initial interval is isomorphic to a product of initial-final intervals.*

Proof. Let $A = [Q_0, A^+]$ be an initial interval. Consider the set of diagonals of Q_0 that belong to A^+. Then we can cut both Q_0 and A^+ along these diagonals. Every piece A_i^- of Q_0 is a smaller Q_0. Every piece A_i^+ of A^+ is a final quadrangulation. By definition of the partial order, the same diagonals belong to every element of A. Cutting along them, we obtain an isomorphism with the product of the intervals $[A_i^-, A_i^+]$, which are initial-final intervals. $\square$

Let us now proceed to a more subtle decomposition.

Let Q be a final triangulation. Let R_0 be the unique region in Q with 0 in its boundary. Removal of R_0 from Q leaves a certain number k of polygons, that will be called the *blocks*. Let us call k the *width* of Q. The width can be 0 only if Q is reduced to R_0. Otherwise it lies between 1 and m.

Proposition 11. *Every initial-final interval A is isomorphic to a product of initial-final intervals of width 1.*

Proof. The main idea is that the flips downwards from Q happen completely independently in distinct blocks. Let us now give a detailed argument for this independence.

Let us consider the set D of diagonals that go from vertex 0 to one of the vertices of the region R_0, except the two vertices that are neighbours of 0. These diagonals are inside the region R_0 and therefore do not belong to Q. The ends of these diagonals receive every label from 0 to $m - 2$ exactly once in clockwise order, and therefore never get the label $m - 1$. Note that the ends of the initial diagonal are labelled by $m - 1$.

We now show by induction downward from Q that, for every $Q' \in [Q_0, Q]$, the diagonals in D do not belong to Q' and do not cross any diagonal of Q'. This is clear for Q. The effect of a down flip is to replace a diagonal by an initial diagonal. This initial diagonal is not in D and does not cross any diagonal in Q, because the labelling of its end is $m - 1$.

We can therefore cut along the diagonals in D. Replacement of every block but a fixed one by a trivial block gives a map to an initial-final interval of width 1. Taking the product of all these maps, we obtain the desired isomorphism.

$\square$

Proposition 12. *The number of elements covered by the maximum in an initial-final interval $[Q_0, Q]$ is the width of Q.*

Proof. In the proof of Lemma 1, it was shown that, in every polygon of Q minus R_0, there is exactly one initial diagonal d that crosses just one diagonal d' of Q. The set of diagonals of Q that can be flipped down in the poset is exactly the collection of these d', and their number is therefore the width of Q. $\square$

In particular, initial-final intervals of width 1 have a maximum that covers a unique element.

Corollary 13. *The intervals in $\mathcal{P}_{m,n}$ can be iteratively built by either adding a maximum to a smaller one or by taking a product of several smaller ones.*

Proof. We proceed by induction on both the height (difference of ranks) of the intervals and the size of the ambient polygon. Every interval is a product of initial-final ones. Every initial-final interval is a product of initial-final intervals of width 1. In both cases, if the product is not reduced to one element, the factors live in strictly smaller polygons.

Every initial-final interval of width 1 is obtained by adding a maximum to an interval, which has smaller height. $\square$

Let us define a *forest poset* as a poset where every element is covered by at most one element. These posets can be iteratively built from the empty poset by two operations, namely adding a maximum element and taking the disjoint union. Their Hasse diagrams are forests of rooted trees, where the roots are the maximal elements.

Proposition 14. *Every interval in $\mathcal{P}_{m,n}$ is a distributive lattice, isomorphic to the lattice of order ideals of a forest poset.*

Proof. By the results above, the intervals in $\mathcal{P}_{m,n}$ can be obtained from smaller intervals using two operations, namely adding a maximum and taking a product. These two operations correspond to adding a maximum or taking the disjoint union, applied on the poset of join-irreducible elements. The statement follows by induction. $\square$

For $m = 1$, this property of intervals was already obtained in [CSS14, § 2].

Corollary 15. *The Möbius function on a poset $\mathcal{P}_{m,n}$ takes its values in $\{-1, 0, 1\}$.*

Proof. This is a well-known property of distributive lattices. $\square$

References

[AB03]   J.-C. Aval and N. Bergeron. Catalan paths and quasi-symmetric functions. *Proc. Amer. Math. Soc.*, 131(4):1053–1062, 2003.

[ABB04] J.-C. Aval, F. Bergeron, and N. Bergeron. Ideals of quasi-symmetric functions and super-covariant polynomials for S_n. *Adv. Math.*, 181(2):353–367, 2004.

[Ava07] J.-C. Aval. Ideals and quotients of B-quasisymmetric polynomials. *Sém. Lothar. Combin.*, 54:Art. B54d, 13, 2005/07.

[BH08]   P. Baumann and C. Hohlweg. A Solomon descent theory for the wreath products $G \wr S_n$. *Trans. Amer. Math. Soc.*, 360(3):1475–1538, 2008.

[BPR12] F. Bergeron and L.-F. Préville-Ratelle. Higher trivariate diagonal harmonics via generalized Tamari posets. *J. Comb.*, 3(3):317–341, 2012.

[Cha05] F. Chapoton. Une base symétrique de l'algèbre des coinvariants quasi-symétriques. *Electron. J. Combin.*, 12:Note 16, 7, 2005.

[CSS12] S. A. Csar, R. Sengupta, and W. Suksompong. On a subposet of the Tamari lattice. In *24th International Conference on Formal Power Series and Algebraic Combinatorics (FPSAC 2012)*, DMTCS Proceedings, pp. 563–574, 2012.

[CSS14] S. A. Csar, R. Sengupta, and W. Suksompong. On a subposet of the Tamari lattice. *Order*, 31(3):337–363, 2014.

[Ges84] Ira M. Gessel. Multipartite *P*-partitions and inner products of skew Schur functions. In *Combinatorics and algebra (Boulder, Colo., 1983)*, Contemp. Math., vol. 34, pp. 289–317. Amer. Math. Soc., Providence, RI, 1984.

[NT10]   Jean-Christophe Novelli and Jean-Yves Thibon. Free quasi-symmetric functions and descent algebras for wreath products, and noncommutative multi-symmetric functions. *Discrete Math.*, 310(24):3584–3606, 2010.

[Pal03] J. M. Pallo. Right-arm rotation distance between binary trees. *Inform. Process. Lett.*, 87(4):173–177, 2003.

[Poi98] S. Poirier. Cycle type and descent set in wreath products. In *Proceedings of the 7th Conference on Formal Power Series and Algebraic Combinatorics (Noisy-le-Grand, 1995)*, vol. 180, pp. 315–343, 1998.

[STW15] C. Stump, H. Thomas, and N. Williams. Cataland: Why the Fuss? arχiv:1503.00710v1.

Institut de Recherche Mathématique Avancée, CNRS UMR 7501, Université de Strasbourg, F-67084 Strasbourg Cedex, France
  E-mail address: chapoton@unistra.fr

LaBRI, Université de Bordeaux, 351, cours de la Libération, 33405 Talence Cedex, France
  E-mail address: jean-christophe.aval@labri.fr