PILES OF CUBES, MONOTONE PATH POLYTOPES AND HYPERPLANE ARRANGEMENTS

CHRISTOS A. ATHANASIADIS

Abstract. Monotone path polytopes arise as a special case of the construction of fiber polytopes, introduced by Billera and Sturmfels. A simple example is provided by the permutahedron, which is a monotone path polytope of the standard unit cube. The permutahedron is the zonotope polar to the braid arrangement. We show how the zonotopes polar to the cones of certain deformations of the braid arrangement can be realized as monotone path polytopes. The construction is an extension of that of the permutahedron and yields interesting connections between enumerative combinatorics of hyperplane arrangements and geometry of monotone path polytopes.

1. Introduction

Fiber polytopes were introduced by Billera and Sturmfels [5] and were further studied in [6]. The fiber polytope $\Sigma(P,Q)$ is a polytope naturally associated to any projection of polytopes $\pi : P \longrightarrow Q$. It is defined as

$$\Sigma(P,Q) = \left\{ \frac{1}{\mathrm{vol}(Q)} \int_Q \gamma(x)\, \mathrm{d}x \mid \gamma \text{ is a section of } \pi \right\}. \tag{1}$$

A concise introduction to the theory of fiber polytopes can be found in Chapter 9 of Ziegler's book [14], on which we rely for general background on polytopes.

We are only interested in the special case in which Q is one dimensional, i.e. a line segment. The polytope $\Sigma(P,Q)$ is then called the *monotone path polytope* of P and π [5, §5] [14, §9.2] and is denoted by $\Pi(P,\pi)$. Its dimension is one less than that of P. The sections of π are the paths in P which are *monotone*, that is strictly increasing with respect to the function π, and *maximal*, i.e. they project to Q. The vertices of $\Pi(P,\pi)$ correspond to certain maximal monotone edge paths on P, called π-*coherent*.

A nice example of a monotone path polytope is the permutahedron $\Pi_{d-1} \subseteq \mathbb{R}^d$. The permutahedron Π_{d-1} is a classical geometric object. It is the convex hull of all vectors in $\mathbb{R}^d$ obtained by permuting the coordinates of the vector $(1,2,\ldots,d)$. It was constructed as a monotone path polytope of the standard unit cube $C_d = [0,1]^d$ by Billera and Sturmfels [5, Example 5.4] (see also [6, Thm. 4.3]). A detailed

Supported by a postdoctoral fellowship from the Mathematical Sciences Research Institute, Berkeley, California. Research at MSRI is supported in part by NSF grant DMS-9022140.

description of the construction can be found in [14, Example 9.8]. The projection used is $\pi : C_d \longrightarrow [0, d]$ with

$$\pi(x) = x_1 + x_2 + \cdots + x_d. \tag{2}$$

The monotone edge paths from the origin to the vertex $(1, 1, \ldots, 1)$ are all π-coherent. They correspond naturally to permutations of a d-element set and give rise to the vertices of Π_{d-1}. This is illustrated in Figure 1 for $d = 3$.

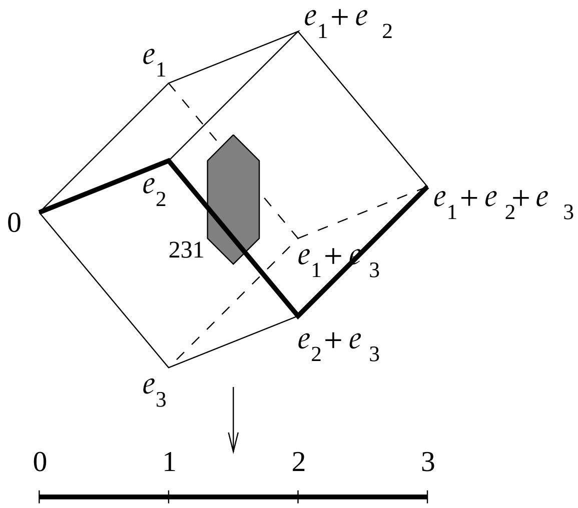

FIGURE 1. The projection π and the permutahedron

An important property of Π_{d-1} is that it is a zonotope. The associated hyperplane arrangement [14, §7.3] is the *braid arrangement* $\mathcal{A}_d$. It consists of the hyperplanes in $\mathbb{R}^d$ of the form $x_i = x_j$ for $1 \le i < j \le d$, i.e. the reflecting hyperplanes of the Coxeter group of type A_{d-1}. Our motivation comes from recent work on *deformations* of $\mathcal{A}_d$, initiated by Stanley [11]. A deformation of $\mathcal{A}_d$ is an affine arrangement which has each of its hyperplanes parallel to some hyperplane $x_i = x_j$ of $\mathcal{A}_d$. We will be interested in the deformations of the form

$$x_i - x_j = -\lambda_i + 1, \ldots, -1, 0, 1, \ldots, \lambda_j - 1 \quad \text{for} \quad 1 \le i < j \le d, \tag{3}$$

where $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_d)$ is a composition of a positive integer n. We denote the arrangement (3) by $\mathcal{A}_d(\lambda)$. For $\lambda_1 = \cdots = \lambda_d = a + 1$ it specializes to the *extended Catalan arrangement* (see for example [11, §2]) and is denoted by $\mathcal{A}_d^{[0,a]}$. This reduces to the braid arrangement $\mathcal{A}_d$ for $a = 0$ and is simply called a *Catalan arrangement* for $a = 1$. Figure 2 shows the Catalan arrangement for $d = 3$, intersected with the hyperplane $x_1 + x_2 + x_3 = 0$.

The combinatorics of deformations of $\mathcal{A}_d$ is very rich and has a growing literature [11, 12, 3, 10]. The zonotopes associated to the corresponding homogenized linear

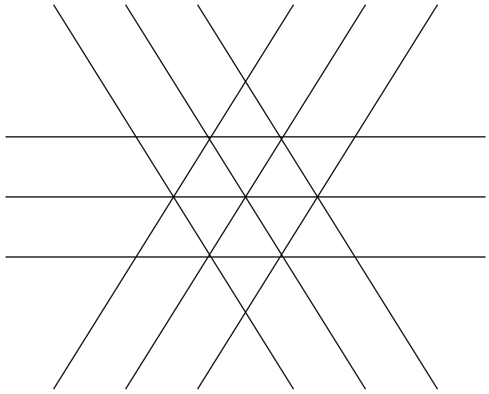

FIGURE 2. The Catalan arrangement for $d = 3$

arrangements, or cones, haven't been studied much as polytopes. They are generalizations of the permutahedron that carry the same interesting combinatorial structure as the arrangements. Our objective is to construct the zonotopes associated to the cones of the arrangements $\mathcal{A}_d(\lambda)$, with their Minkowski summands suitably rescaled, as monotone path polytopes.

The construction generalizes that of the permutahedron. In particular, the cube C_d is replaced by Ziegler's *lifted pile of cubes* $\widetilde{\mathcal{P}}_{d+1}(\lambda)$ [14, §5.1], which we define formally in the next section. The projection π is given by the formula (2). The corresponding monotone path polytope $\Pi_d(\lambda)$ is a zonotope, as described above, with its upper part truncated. The maximal monotone edge paths in the lower faces of $\widetilde{\mathcal{P}}_{d+1}(\lambda)$ are in bijection with the lattice paths in $\mathbb{R}^d$ from the origin to the point λ, having unit coordinate steps. In general, most of these paths are *not* coherent. The coherent ones correspond to the lower vertices of the zonotope and hence to the regions of $\mathcal{A}_d(\lambda)$. This will enable us to enumerate them (see Corollary 4.3) by an easy application of the "finite field method" of [1, 2]. It follows that their number is small compared to the total number of lower monotone edge paths.

The paper is organized as follows: In Section 2 we give the necessary definitions and state our main results about $\Pi_d(\lambda)$. Section 3 contains a proof of the main theorem. In Section 4 we draw all enumerative consequences of the main theorem by analyzing the combinatorics of $\mathcal{A}_d(\lambda)$. We also characterize the coherent paths in the lower faces of $\widetilde{\mathcal{P}}_{d+1}(\lambda)$ directly from the definition. We close with some remarks in Section 5.

2. Basic definitions and results

We begin with some notation and terminology. We denote by $e_1, e_2, \ldots, e_d$ the unit coordinate vectors in $\mathbb{R}^d$. We write $[a, b]$ for the line segment joining two points

$a, b \in \mathbb{R}^d$. The *canonical projection* $\mathbb{R}^{d+1} \longrightarrow \mathbb{R}^d$ is the map which simply forgets the last coordinate. Lastly, we refer to the union of the lower faces [14, §5.1], or upper faces, of a polytope P as its *lower part*, or *upper part* respectively.

We first give the definition of the polytope $\widetilde{\mathcal{P}}_{d+1}(\lambda)$, as promised in the introduction. Let $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_d)$ be a composition of a positive integer n. In other words, the λ_i are positive integers which sum to n. The *pile of cubes* $\mathcal{P}_d(\lambda)$, corresponding to λ, is the polytopal complex formed by all unit cubes with integer vertices in the *d-box*

$$B(\lambda) = \left\{ (x_1, \dots, x_d) \in \mathbb{R}^d \mid 0 \le x_i \le \lambda_i \text{ for } 1 \le i \le d \right\}.$$

In particular, the vertex set of $\mathcal{P}_d(\lambda)$ is

$$\operatorname{vert}(\mathcal{P}_d(\lambda)) = B(\lambda) \cap \mathbb{Z}^d.$$

Let $f : \mathbb{R}^d \longrightarrow \mathbb{R}$ be of the form

$$f(x) = f_1(x_1) + \cdots + f_d(x_d) \tag{4}$$

where the f_i's are strictly convex functions of one variable. For the purpose of our results we will use the canonical choice

$$f(x) = x_1^2 + x_2^2 + \cdots + x_d^2.$$

We take this as the definition of f from now and on. The *lifted pile of cubes* corresponding to λ is the polytope in $\mathbb{R}^{d+1}$

$$\widetilde{\mathcal{P}}_{d+1}(\lambda) = \operatorname{conv}\left\{ (x, f(x)) \mid x \in \operatorname{vert}(\mathcal{P}_d(\lambda)) \right\}.$$

The pile of cubes $\mathcal{P}_d(\lambda)$ and its lift $\widetilde{\mathcal{P}}_{d+1}(\lambda)$ are discussed in [14, Example 5.4]. For suitable λ they provide examples of polytopal complexes, respectively polytopes, which are not extendably shellable [14, Ch. 8]. Figure 3 shows $\mathcal{P}_2(4,3)$ and its lift $\widetilde{\mathcal{P}}_3(4,3)$.

The lower faces of $\widetilde{\mathcal{P}}_{d+1}(\lambda)$ form a polytopal complex combinatorially equivalent to $\mathcal{P}_d(\lambda)$. Because of the form (4), these faces are parallelepipeds, as Figure 3 suggests. By construction, they map to the corresponding faces of $\mathcal{P}_d(\lambda)$ under the canonical projection $\mathbb{R}^{d+1} \longrightarrow \mathbb{R}^d$. The unique upper facet, defined by the hyperplane

$$\lambda_1 x_1 + \lambda_2 x_2 + \cdots + \lambda_d x_d = x_{d+1}, \tag{5}$$

is a parallelepiped which projects to $B(\lambda)$.

Let $\Pi_d(\lambda)$ be the monotone path polytope of the projection

$$\pi : \widetilde{\mathcal{P}}_{d+1}(\lambda) \longrightarrow [0, n]$$

given by the formula (2). Thus $\pi(x)$ is the sum of the first d coordinates of $x \in \mathbb{R}^{d+1}$. We have $\dim \Pi_d(\lambda) = \dim \widetilde{\mathcal{P}}_{d+1}(\lambda) - 1 = d$, the number of parts of λ except if $\lambda_i = 1$ for all i. In this case $\widetilde{\mathcal{P}}_{d+1}(\lambda)$ is affinely isomorphic to the d-cube C_d, so $\Pi_d(\lambda)$ is affinely isomorphic to the permutahedron Π_{d-1} and $\dim \Pi_d(\lambda) = d - 1$.

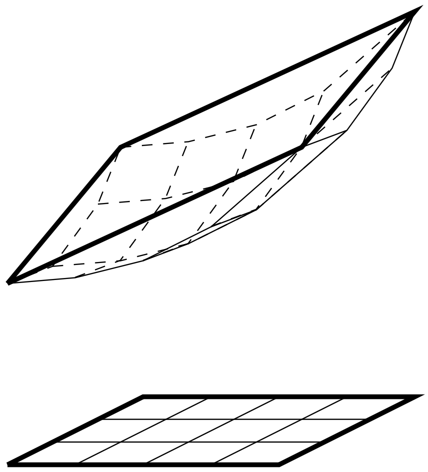

FIGURE 3. $\mathcal{P}_2(4,3)$ and $\widetilde{\mathcal{P}}_3(4,3)$

A *truncation* of a polytope P in $\mathbb{R}^r$ is its intersection with a closed halfspace $a \cdot x \leq 0$, where $a \in \mathbb{R}^r$. We call such a truncation an *upper truncation* if it includes the lower part of P but does not intersect the upper part in the region $a \cdot x < 0$. In particular, an upper truncation has a unique upper facet. It rarely happens that a polytope P has an upper truncation. However, if P is a zonotope with a sufficiently large Minkowski summand in the direction of the last coordinate, as in the situation of the next theorem, then clearly, upper truncations exist.

Theorem 2.1. *Let $Z_d(\lambda)$ be the zonotope*

$$[0, s\, e_{d+1}] + \sum_{1 \leq i < j \leq d} \sum_{k=0}^{\lambda_i - 1} \sum_{l=0}^{\lambda_j - 1} [0, e_j - e_i + (2l - 2k)\, e_{d+1}]\,,$$

where s is a sufficiently large positive number. The polytope $\Pi_d(\lambda)$ is an upper truncation of a translate of $\frac{1}{n} Z_d(\lambda)$, cut by the hyperplane (5).

Note that the hyperplane arrangement associated to $Z_d(\lambda)$ is projectively equivalent to the *cone* of $\mathcal{A}_d(\lambda)$, denoted $\mathbf{c}\mathcal{A}_d(\lambda)$. This is the linear arrangement in $\mathbb{R}^{d+1}$ obtained by homogenizing each hyperplane $x_i - x_j = s$ of $\mathcal{A}_d(\lambda)$ to $x_i - x_j = s x_{d+1}$ and adding the hyperplane $x_{d+1} = 0$. Hence $Z_d(\lambda)$ is combinatorially equivalent to the zonotope $Z[\mathbf{c}\mathcal{A}_d(\lambda)]$, associated to this cone, which we call the *polar zonotope*. The faces of $Z[\mathbf{c}\mathcal{A}_d(\lambda)]$ are in inclusion-reversing bijection with the faces of $\mathbf{c}\mathcal{A}_d(\lambda)$ [14, Cor. 7.17]. Its lower faces define a zonotopal complex $Z[\mathcal{A}_d(\lambda)]$ which we call the *polar zonotopal complex*, or simply the *polar complex*. The faces of $Z[\mathcal{A}_d(\lambda)]$ are in inclusion-reversing bijection with those of $\mathcal{A}_d(\lambda)$. Figure 4 shows the complex polar to the Catalan arrangement of Figure 2, projected canonically on $\mathbb{R}^2$.

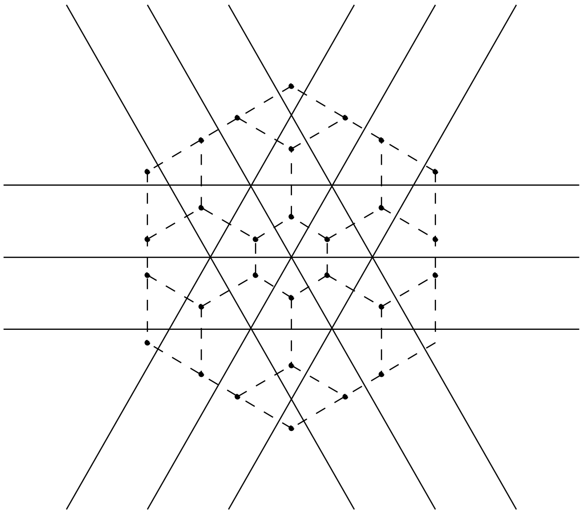

FIGURE 4. The polar complex

We now focus on the the lower part, which is the interesting part of $\Pi_d(\lambda)$. Theorem 2.1 implies the following corollary.

Corollary 2.2. *The complex of all lower faces of* $\Pi_d(\lambda)$ *is combinatorially equivalent to the polar complex* $Z[\mathcal{A}_d(\lambda)]$ *of the arrangement* $\mathcal{A}_d(\lambda)$. *In particular, the lower vertices of* $\Pi_d(\lambda)$ *are in bijection with the regions of* $\mathcal{A}_d(\lambda)$. □

Hence, for $\lambda = (2,2,2)$, $\Pi_d(\lambda)$ looks combinatorially like the complex of Figure 4, from a point far below on the x_3 axis.

It follows from Corollary 2.2 that the π-coherent monotone edge paths in the lower part of $\widetilde{\mathcal{P}}_{d+1}(\lambda)$ are in bijection with the regions of $\mathcal{A}_d(\lambda)$. To characterize them, we introduce some more terminology. We call the maximal monotone edge paths (with respect to π) in $\mathcal{P}_d(\lambda)$ simply *monotone λ-paths.* They are the lattice paths in $\mathbb{R}^d$ from the origin to the point λ having n unit coordinate steps. They correspond naturally to permutations of the multiset M_λ, which contains i with multiplicity λ_i for $1 \le i \le d$. We call these permutations *λ-permutations* and denote by p^w the monotone λ-path which corresponds to w. For example, if $\lambda = (2,2,2)$ and $w = 121323$ then the path p^w has successive steps $e_1, e_2, e_1, e_3, e_2, e_3$. The number of λ-permutations is the multinomial coefficient $\binom{n}{\lambda_1, \ldots, \lambda_d}$.

The maximal monotone edge paths in the lower part of $\widetilde{\mathcal{P}}_{d+1}(\lambda)$ are in bijection with the monotone λ-paths, via the canonical projection $\mathbb{R}^{d+1} \longrightarrow \mathbb{R}^d$. We refer to these paths in $\widetilde{\mathcal{P}}_{d+1}(\lambda)$ as *lifted monotone λ-paths* and use the notation γ^w.

A λ-permutation $w = w_1 w_2 \cdots w_n$ is *nesting* if there exist indices $1 \le i < j < k < l \le n$ such that $w_i = w_l$, $w_j = w_k$ and $w_m \ne w_i$ for all m with $j \le m \le k$. Otherwise

w is *non-nesting*. Equivalently, w is non-nesting if the following linear system in the variables $y_1, y_2, \dots, y_n$ is feasible:

(*i*) $y_1 < y_2 < \cdots < y_n$,

(*ii*) $y_k = y_j + 1$ if $1 \leq j < k \leq n$, $w_k = w_j$ and $w_m \neq w_j$ for $j < m < k$.

For example, 12121 is non-nesting but 12211 is not.

Theorem 2.3. *The lifted monotone λ-path γ^w is π-coherent if and only if w is non-nesting. The number of such paths, or equivalently the number of lower vertices of $\Pi_d(\lambda)$, or regions of $\mathcal{A}_d(\lambda)$ is $\frac{n!}{(n-d+1)!}$. In particular, this quantity depends only on n and d.*

It follows that the fraction of monotone λ-paths with π-coherent lifts is

$$\frac{\lambda_1!\lambda_2!\cdots\lambda_d!}{(n-d+1)!}.$$

This result gives concrete examples of projections for which the fraction of coherent monotone edge paths is explicitly shown to be small. For instance, in the case $\lambda_1 = \cdots = \lambda_d = 2$ this quantity becomes $2^d/(d+1)!$. The only other fiber polytope situation we are aware of, in which it has been shown that the number of coherent subdivisions is asymptotically negligible compared to the number of all subdivisions, appears in [7] in the context of triangulations of cyclic d-polytopes with $d+4$ vertices.

3. Proof of the main theorem

We now prove Theorem 2.1.

Proof of Theorem 2.1: We use directly the defining formula (1). Recall that the vertices of $\Pi_d(\lambda)$ correspond to certain maximal monotone edge paths γ in $\widetilde{\mathcal{P}}_{d+1}(\lambda)$. Since $\pi(x)$ depends only on the first d coordinates of x, we only need to consider such paths either in the lower or upper part of $\widetilde{\mathcal{P}}_{d+1}(\lambda)$ and take the convex hull of the right hand side of (1). Let

$$I_\gamma := \frac{1}{\mathrm{vol}(Q)} \int_{[0,n]} \gamma(x)\,\mathrm{d}x = \frac{1}{n} \sum_{k=1}^{n} \frac{\gamma(k-1)+\gamma(k)}{2}.$$

We first consider the lower part of $\widetilde{\mathcal{P}}_{d+1}(\lambda)$, so the relevant paths are the lifted monotone λ-paths γ^w. Instead of calculating I_{γ^w} explicitly, we note that any λ-permutation w can be obtained from $w_0 = 1\cdots1\cdots d\cdots d$ by successively swapping adjacent entries ij with $i < j$. Let $w = w_1 w_2 \cdots w_n$, where $w_k = i < j = w_{k+1}$ and let w' be obtained from w by swapping w_k and w_{k+1}. If $v := \gamma^w(k-1) = \gamma^{w'}(k-1) = (m_1, m_2, \dots, m_d)$, then

$$I_{\gamma^{w'}} - I_{\gamma^w} = \tfrac{1}{n}\left(\gamma^{w'}(k) - \gamma^w(k)\right) =$$

$$\tfrac{1}{n}\,(v + e_j + f(v+e_j)\,e_{d+1} - v - e_i - f(v+e_i)\,e_{d+1}) =$$

$$\tfrac{1}{n}\,(e_j - e_i + (2m_j - 2m_i)\,e_{d+1})\,.$$

It follows easily that the points I_{γ^w} lie on $\frac{1}{n}Z_d(\lambda) + I_{\gamma^{w_0}}$ and include all its lower vertices.

Finally suppose that γ lies in the upper facet of $\widetilde{\mathcal{P}}_{d+1}(\lambda)$, which is defined by the hyperplane (5). This upper facet projects to the d-box $B(\lambda)$ under the canonical projection $\mathbb{R}^{d+1} \longrightarrow \mathbb{R}^d$. Each of the monotone edge paths γ projects to a monotone λ-path with corresponding λ-permutation of the form $w = \sigma_1 \cdots \sigma_1 \cdots \sigma_d \cdots \sigma_d$, where $\sigma_1 \cdots \sigma_d$ is a permutation of $\{1, 2, \ldots, d\}$. Hence $I_\gamma = I_{\gamma^w} + te_{d+1}$ for some $t \geq 0$. Since γ lies on (5), so does I_γ. It follows that the upper part of $\Pi_d(\lambda)$ also consists of a single facet, defined by (5). It projects to the same zonotope on $\mathbb{R}^d$, which is combinatorially equivalent to the permutahedron Π_{d-1}, as the lower part. This completes the proof. $\square$

4. The coherent paths and enumeration

In this section we characterize and enumerate the π-coherent lifted monotone λ-paths and thus prove Theorem 2.3. We proceed directly from the definition of a coherent path and extend our argument to get another proof of Corollary 2.2. For the enumeration we use the finite field method of [1, 2].

We first recall what it means for a monotone edge path to be coherent. Let P be a polytope in $\mathbb{R}^r$, $\pi : \mathbb{R}^r \longrightarrow \mathbb{R}$ an affine function, $Q = \pi(P) \subseteq \mathbb{R}$ and p a maximal monotone edge path in P. The path p is said to be π-*coherent* if there exists a *generic* linear functional $c \in (\mathbb{R}^r)^*$ such that $\pi^c : \mathbb{R}^r \longrightarrow \mathbb{R}^2$ with

$$\pi^c(x) = (\pi(x), c(x))$$

maps p to the path of lower edges of the polygon $Q^c := \pi^c(P)$. More generally [14, Definition 9.2], a set of faces $\mathcal{F}$ of P is said to form a π-*coherent subdivision* of Q if, for some $c \in (\mathbb{R}^r)^*$, $\mathcal{F}$ is the set $\mathcal{F}^c$ of faces $(\pi^c)^{-1}(G)$ of P, where G runs through the lower faces of Q^c. Note that $\mathcal{F}^c$ induces the subdivision $\{\pi(F) : F \in \mathcal{F}^c\}$ of Q. This is also the subdivision obtained by projecting canonically the lower edges of Q^c on Q. The set $\mathcal{F}^c$ consists of faces of P, rather than Q.

For the rest of this section let $r = d+1$, $P = \widetilde{\mathcal{P}}_{d+1}(\lambda)$, $Q = [0, n]$ and π as in Section 2.

Proposition 4.1. *The lifted monotone λ-path γ^w is π-coherent if and only if w is non-nesting. The non-nesting λ-permutations are in bijection with the regions of the arrangement $\mathcal{A}_d(\lambda)$.*

Proof. Let $c \in (\mathbb{R}^{d+1})^*$ be generic and given by

$$c(x) = a_1 x_1 + a_2 x_2 + \cdots + a_{d+1} x_{d+1}. \tag{6}$$

We need to describe the π-coherent path defined by c, directly in terms of c. Since we want this path to lie in the lower part of $\widetilde{\mathcal{P}}_{d+1}(\lambda)$, we can assume that a_{d+1} is positive, say $a_{d+1} = 1/2$ by rescaling.

Let $w = w_1 w_2 \cdots w_n$ be a λ-permutation and p^w, γ^w as in Section 2, with $p^w = (v_0, v_1, \ldots, v_n)$ and $\gamma^w = (\tilde{v}_0, \tilde{v}_1, \ldots, \tilde{v}_n)$. Thus $v_k = v_{k-1} + e_i$ if $w_k = i$ and $\tilde{v}_k = (v_k, f(v_k))$ for all k. We denote by $\Delta^c(w)$ the sequence of length n whose kth term records the difference

$$c(\tilde{v}_k) - c(\tilde{v}_{k-1})$$

of c at the kth edge of γ^w. Let $v_{k-1} = (m_1, \ldots, m_d)$ and $w_k = i$. Then

$$c(\tilde{v}_k) - c(\tilde{v}_{k-1}) = a_i + a_{d+1} \left(f(v_{k-1} + e_i) - f(v_{k-1}) \right) =$$

$$a_i + (1/2)\,(2m_i + 1) = a'_i + m_i,$$

where $a'_i = a_i + 1/2$ for $1 \leq i \leq d$ and the possible values of m_i are the integers satisfying $0 \leq m_i \leq \lambda_i - 1$. It follows that $\Delta^c(w)$ is the permutation of the numbers

$$a'_i + m, \quad 0 \leq m \leq \lambda_i - 1, \ 1 \leq i \leq d \tag{7}$$

which places $a'_i, a'_i + 1, \ldots, a'_i + \lambda_i - 1$, in this order, in the λ_i positions occupied by i in w. The numbers (7) are distinct because c is assumed to be generic.

Let w^c be the λ-permutation which corresponds to the π-coherent path defined by c. If π^c maps γ^w to the lower edge path of the polygon Q^c, then $\Delta^c(w)$ is strictly increasing, by the convexity of this path. It follows that w^c is the unique λ-permutation w for which the permutation $\Delta^c(w)$ of the numbers (7) is strictly increasing. Let $c' = (a'_1, a'_2, \ldots, a'_d) \in \mathbb{R}^d$. For a fixed λ-permutation w, if nonempty, the set of all c' such that $w^c = w$ forms exactly a region of $\mathcal{A}_d(\lambda)$. Indeed, deciding which of the two numbers $a'_i + m$ and $a'_j + l$ in (7) is smaller amounts to choosing a side of the hyperplane $x_i - x_j = l - m$ of $\mathcal{A}_d(\lambda)$. Finally, $w = w^c$ for some c if and only if the condition that $\Delta^c(w)$ is strictly increasing does not impose inequalities of the form

$$d_j < d_i < d_i + 1 < d_j + 1.$$

By definition, this means that w is non-nesting. $\square$

We now compute the number of regions of $\mathcal{A}_d(\lambda)$. We assume familiarity with the *characteristic polynomial* $\chi(\mathcal{A}, q)$ of $\mathcal{A}$ [9, §2.3], a fundamental combinatorial invariant of $\mathcal{A}$, and use Zaslavsky's theorem [13] for the number of regions of $\mathcal{A}$. The characteristic polynomial of $\mathcal{A}_d(\lambda)$ can be computed by an easy application of the finite field method of [1] [2, Part II]. This method reduces the computation to a simple counting problem in a vector space over a finite field. The argument in the following proposition is similar to the one given in [1, Thm. 5.1] for the special case of the extended Catalan arrangements.

Proposition 4.2. *We have*

$$\chi(\mathcal{A}_d(\lambda), q) = q \prod_{j=n-d+1}^{n-1} (q-j).$$

Proof. Let q be a large prime number and let $\mathbb{F}_q$ denote the finite field of integers mod q. Theorem 2.2 in [1] (also [2, Thm. 5.2.1]) implies that $\chi(\mathcal{A}_d(\lambda), q)$ counts the number of d-tuples $(x_1, x_2, \ldots, x_d) \in \mathbb{F}_q^d$ satisfying

$$x_i - x_j \neq -\lambda_i + 1, \ldots, -1, 0, 1, \ldots, \lambda_j - 1 \ \text{ for } \ 1 \leq i < j \leq d.$$

Equivalently, we want to choose the x_i so that the classes mod q

$$x_i, x_i + 1, \ldots, x_i + \lambda_i - 1 \ \text{ for } \ 1 \leq i \leq d \tag{8}$$

are distinct. To count these d-tuples we first cyclically permute the d strings (8) in $(d-1)!$ ways. Then we distribute $q-n$ indistinguishable boxes in the d spaces between successive strings in $\binom{q-n+d-1}{d-1}$ ways. The $q-n$ boxes distributed stand for the $q-n$ classes in $\mathbb{F}_q$ not of the form (8). Finally we assign to x_1 a specific value in $\mathbb{F}_q$ in q ways, say $x_1 = 0$. The other classes of (8) $x_1 + 1, \ldots$ and the boxes are naturally assigned the values $1, \ldots, q-1 \in \mathbb{F}_q$, according to their cyclic arrangement. The product

$$q\,(d-1)! \binom{q-n+d-1}{d-1}$$

is the expression for $\chi(\mathcal{A}_d(\lambda), q)$ we have claimed. □

Zaslavsky's theorem [13] expresses the number of regions of an arrangement $\mathcal{A}$ in $\mathbb{R}^d$ as $(-1)^d \chi(\mathcal{A}, -1)$ and yields the following corollary.

Corollary 4.3. *The number of regions of $\mathcal{A}_d(\lambda)$ is*

$$\frac{n!}{(n-d+1)!}.$$

In particular, this number depends only on n and d, the sum and number of parts of λ respectively. □

Figure 5 shows the arrangement $\mathcal{A}_d(\lambda)$ for $\lambda = (3, 2, 1)$, intersected with the hyperplane $x_1 + x_2 + x_3 = 0$. Note that it has 30 regions, as many as the Catalan arrangement of Figure 4, which corresponds to $\lambda = (2, 2, 2)$. Corollary 4.3 specializes to the formula given in [11, §2] for the number of regions of the extended Catalan arrangement $\mathcal{A}_d^{[0,a]}$, since in this case $n = (a+1)d$. Proposition 4.1 and Corollary 4.3 imply Theorem 2.3.

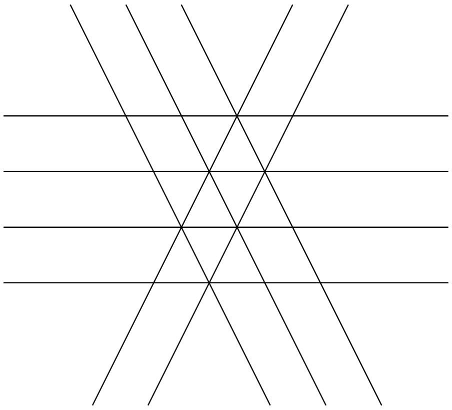

FIGURE 5. The arrangement $\mathcal{A}_3(3,2,1)$

The reasoning in the proof of Proposition 4.1 can be extended to determine the face lattice of $\Pi_d(\lambda)$. Below we give an alternative proof of Corollary 2.2, obtained in this way. By the general theory of fiber polytopes [5, Thm. 2.1] [14, Thm. 9.6], the face lattice of $\Pi_d(\lambda)$ is isomorphic to the poset of π-coherent subdivisions of $Q = [0, n]$, with a minimal element adjoined. In this poset we have $\mathcal{F}_1 \leq \mathcal{F}_2$ if the union of the faces in $\mathcal{F}_1$ is contained in the union of the faces in $\mathcal{F}_2$.

Alternative proof of Corollary 2.2: To get the lower faces of $\Pi_d(\lambda)$ we consider linear functionals $c \in (\mathbb{R}^{d+1})^*$ as in (6), with positive a_{d+1}, but not necessarily generic. Let P^c denote the set of monotone λ-paths whose lifts lie in $\mathcal{F}^c$, so that P^c is a singleton if c is generic. Note that $\mathcal{F}^{c_1} \leq \mathcal{F}^{c_2}$ in the poset of π-coherent subdivisions if and only if $P^{c_1} \subseteq P^{c_2}$. Using the notation in the proof of Proposition 4.1, we have $p^w \in P^c$ if and only if $\Delta^c(w)$ is *weakly* increasing. It follows that $P^{c_1} \subseteq P^{c_2}$ if and only if the faces F_1, F_2 of $\mathcal{A}_d$ which contain c'_1 and c'_2 respectively satisfy $F_1 \geq F_2$ in the face poset of $\mathcal{A}_d$. $\square$

A precise description of the π-coherent lower subdivisions is given in Remark 2 of the next section. The computations of general face numbers of the extended Catalan arrangements [2, Ch. 8] can be translated in terms of the lower part of $\Pi_d(\lambda)$ via Corollary 2.2. The following corollary, for example, follows from [2, Cor. 8.3.2], which computes the face numbers of the Catalan arrangement.

Corollary 4.4. *For $\lambda = (2, 2, \ldots, 2)$ and $1 \leq k \leq d$, the number of lower faces of $\Pi_d(\lambda)$ of dimension $d - k$ is*

$$\sum_{r=k}^{d} (r-1)!\, S(d, r) \binom{r}{k} \binom{r+k}{k-1},$$

where $S(d, k)$ stands for a Stirling number of the second kind. $\square$

5. Remarks

1. A direct bijective proof of Corollary 4.3 is possible and is outlined next.

Alternative proof of Corollary 4.3: Recall that a region of $\mathcal{A}_d(\lambda)$ is defined by a linear ordering $\tau_1 < \tau_2 < \cdots < \tau_n$ of

$$x_i + m, \quad 0 \le m \le \lambda_i - 1, \ 1 \le i \le d$$

which respects the orderings $x_i < x_i + 1 < \cdots < x_i + \lambda_i - 1$ for $1 \le i \le d$ and is such that the λ-permutation obtained by replacing each $x_i + m$ with i is non-nesting. Given such an ordering τ with corresponding permutation $w = w_1 w_2 \cdots w_n$, let $j_1, j_2, \ldots, j_d$ be the positions in which $x_1, x_2, \ldots, x_d$ appear in τ respectively. In other words, j_i is the smallest j such that $w_j = i$. Consider the quotient of the abelian group $\mathbb{Z}^d_{n+1}$ by the cyclic subgroup H generated by $(1, 1, \ldots, 1)$. The map

$$\tau \longrightarrow J_\tau = (j_1, j_2, \ldots, j_d) + H$$

defines a bijection between the regions of $\mathcal{A}_d(\lambda)$ and cosets $J = (j_1, j_2, \ldots, j_d) + H$ of $\mathbb{Z}^d_{n+1}/H$ for which all j_i are mutually distinct. Clearly, the number of such cosets is $n(n-1)\cdots(n-d+2)$. □

2. Let $\pi : P \longrightarrow Q$ be as in the beginning of Section 4. The poset of π-coherent subdivisions of Q is an induced subposet of the poset $\omega(P, Q)$ of all *π-induced subdivisions* of Q [14, Definition 9.1] (see also [4, §3]). In the special case $\dim Q = 1$ we are considering, this is the poset of cellular strings of P, induced by π [4, §1]. The π-induced subdivisions in the lower part $P = \widetilde{\mathcal{P}}_{d+1}(\lambda)$ form an interesting poset, which we denote by $\Omega(\lambda)$. It is the set of all *proper* ordered partitions of the multiset M_λ, partially ordered by refinement. An ordered partition of M_λ is proper if none of its blocks contains repeated elements. The subdivision defined by $\rho = (B_1, B_2, \ldots, B_k) \in \Omega(\lambda)$ has as maximal faces the lifts of the faces

$$\sum_{i=1}^{j-1} e_{B_i} + \sum_{r \in B_j} [0, e_r], \quad 1 \le j \le k$$

of $\mathcal{P}_d(\lambda)$, where $e_B = \sum_{i \in B} e_i$. Figure 6 shows the faces of $\mathcal{P}_2(4,3)$ which correspond to the proper ordered partition $(1, 12, 2, 12, 1)$ of $M_\lambda = \{1,1,1,1,2,2,2\}$, with $\lambda = (4,3)$. The minimal elements of $\Omega(\lambda)$ are the λ-permutations, which give rise to the lifted monotone λ-paths.

The argument in the proof of Corollary 2.2 in Section 4 yields the following description of the subposet of π-coherent lower subdivisions: Let $\rho = (B_1, B_2, \ldots, B_k) \in \Omega(\lambda)$, as before. For each $1 \le i \le d$ replace the entries in the blocks of ρ equal to i, from left to right, with $x_i, x_i + 1, \ldots, x_i + \lambda_i - 1$ respectively, to get $\rho' = (B'_1, B'_2, \ldots, B'_k)$. The π-induced subdivision defined by ρ is coherent if and only if the following linear system in the variables $x_1, x_2, \ldots, x_d$ is feasible:

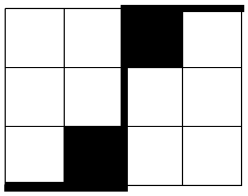

FIGURE 6. The subdivision defined by $(1, 12, 2, 12, 1)$

(*i*) $u = v$ whenever u, v are in the same block B'_j,
(*ii*) $u < v$ whenever $u \in B'_j$, $v \in B'_l$ and $j < l$.

In this case, the set of solutions forms exactly a face of $\mathcal{A}_d(\lambda)$. The subdivision of Figure 6 is not coherent because the system

$$x_1 < x_1 + 1 = x_2 < x_2 + 1 < x_1 + 2 = x_2 + 2 < x_1 + 3$$

has no solution. The only coherent atom of $\Omega(4, 3)$ smaller than this subdivision is $(1, 2, 1, 2, 1, 2, 1)$.

3. There is a notion of "flip" on the set of all maximal monotone edge paths p of P with respect to the affine function π. Two such paths are said to be related by a *flip* if they have a common cover in the poset $\omega(P, Q)$ of π-induced subdivisions of Q. We call the minimum number of flips required to make p coherent the *incoherency* of p. In the case of lifted monotone λ-paths γ^w, the operation of flipping swaps two distinct successive entries ij of w. Therefore the incoherency of γ^w is the minimum number of swappings needed to make w non-nesting. To get a feeling for this number, we compute the maximum incoherency that can occur for $\lambda = (2, 2, \ldots, 2)$.

Proposition 5.1. *Let* $\lambda = (2, 2, \ldots, 2)$. *The maximum incoherency of a path* γ^w *is* $\binom{d}{2}$ *and is attained by the* λ*-permutations of the form* $w_\sigma = \sigma_1\, \sigma_2 \cdots \sigma_d\, \sigma_d \cdots \sigma_2\, \sigma_1$, *where* $\sigma = \sigma_1\, \sigma_2 \cdots \sigma_d$ *is a permutation of* $\{1, \ldots, d\}$.

Proof. Every flip can reduce the number of "nestings" of w by at most one and w^σ has the maximum number $\binom{d}{2}$ of nestings. On the other hand, any $w = w_1 w_2 \cdots w_{2d}$ can be turned into some w_σ by at most $\binom{d}{2}$ flips in the following way: Suppose that $w_1, \ldots, w_i$ are mutually distinct but $w_{i+1} = w_k$ for some $1 \le k \le i$. Let j be the smallest index with $j > i$ for which $w_j \neq w_1, \ldots, w_i$. Move w_j to the left until it occupies position $i + 1$. Continue in the same way until the first d entries of the permutation are all distinct, say $\sigma_1, \sigma_2, \ldots, \sigma_d$. Now continue flipping in the second half of the permutation, in an obvious way, to get w_σ for $\sigma = \sigma_1\, \sigma_2 \cdots \sigma_d$. Any two distinct integers i, j are flipped at most once during the whole process, so the total number of flips is at most $\binom{d}{2}$, as desired. $\square$

4. The extended Catalan arrangements form one of the two families of deformations of $\mathcal{A}_d$ for which explicit formulas have been obtained for the number of faces of any given dimension [2, Ch. 8]. The other family consists of the *extended Shi arrangements*

$$x_i - x_j = -a+1, -a+2, \dots, a \quad \text{for} \quad 1 \le i < j \le d,$$

where $a \geq 1$ is an integer. The number of faces in each dimension was shown to have a surprisingly simple combinatorial interpretation [2, Thm. 8.2.1]. We don't know if a fiber polytope construction for the polar complexes exists in this case.

5. The cones of the extended Catalan arrangements $\mathcal{A}_d^{[0,a]}$ were shown to be *inductively free* [9, Ch. 4] by Edelman and Reiner (see the proof of Theorem 3.2 in [8]). Other classes of deformations of $\mathcal{A}_d$, including the extended Shi arrangements, were shown to be inductively free in [3]. Proposition 4.2 suggests that the same is true for the arrangements $\mathcal{A}_d(\lambda)$.

Acknowledgement. I am grateful to Lou Billera and Bernd Sturmfels for helpful discussions.

References

[1] C. A. Athanasiadis, Characteristic polynomials of subspace arrangements and finite fields, *Advances in Math.* **122** (1996), 193–233.
[2] C. A. Athanasiadis, *Algebraic combinatorics of graph spectra, subspace arrangements and Tutte polynomials*, Ph.D. thesis, MIT, 1996.
[3] C. A. Athanasiadis, On free deformations of the braid arrangement, Preprint 526/1996, TU-Berlin, July 1996, 17 pages.
[4] L. J. Billera, M. M. Kapranov and B. Sturmfels, Cellular strings on polytopes, *Proc. Amer. Math. Soc.* **122** (1994), 549–555.
[5] L. J. Billera and B. Sturmfels, Fiber polytopes, *Annals of Math.* **135** (1992), 527–549.
[6] L. J. Billera and B. Sturmfels, Iterated fiber polytopes, *Mathematika* **41** (1994), 348–363.
[7] J. A. de Loera, S. Hoşten, F. Santos and B. Sturmfels, The polytope of all triangulations of a point configuration, *Doc. Math.* **1** (1996), 103–119.
[8] P. H. Edelman and V. Reiner, Free arrangements and rhombic tilings, *Discrete Comput. Geom.* **15** (1996), 307–340.
[9] P. Orlik and H. Terao, *Arrangements of Hyperplanes*, Grundlehren 300, Springer-Verlag, New York, NY, 1992.
[10] A. Postnikov and R. Stanley, Deformations of Coxeter hyperplane arrangements, in preparation.
[11] R. Stanley, Hyperplane arrangements, interval orders and trees, *Proc. Nat. Acad. Sci.* **93** (1996), 2620–2625.
[12] R. Stanley, Hyperplane arrangements, parking functions and tree inversions, in *Festschrift in Honor of Gian-Carlo Rota*, Birkhäuser, Boston/Basel/Berlin, to appear.
[13] T. Zaslavsky, Facing up to arrangements: face-count formulas for partitions of space by hyperplanes, *Mem. Amer. Math. Soc.* vol. 1, no. 154, (1975).

[14] G. M. Ziegler, *Lectures on Polytopes*, Graduate Texts in Mathematics **152** Springer-Verlag, New York, 1995.

Christos A. Athanasiadis, Mathematical Sciences Research Institute, 1000 Centennial Drive, Berkeley, CA 94720

E-mail address: athana@msri.org