FROM POLYTOPES TO ENUMERATION

ED SWARTZ

1. Preview

What should the d-dimensional analogues of polygons be? The short answer is "convex d-polytopes". We may not know what these are yet, but that does not stop us from asking a number of questions about them.

Example 1.1. *The cube and the octahedron. The cube has 8 vertices, 12 edges and 6 faces, while the octahedron has 6 vertices, 12 edges and 8 faces. Is this reversal in the number of faces in each dimension common?*

In the problem session you saw that v-e+f=2 for 3-polytopes. Is something similar true in higher dimensions? Are there other restrictions on the number of faces of each dimension? What is the most number of faces in each dimension if we fix the number of vertices and dimension? The least? What if we fix other numbers of faces in random dimensions?

Notice that all of these questions involve counting faces of various dimensions. Another situation where these types of questions occur comes is hyperplane arrangements. A *hyperplane* of $\mathbb{R}^d$ is any set of the form $l_{\mathbf{v}}^{-1}(c) = \{\mathbf{x} : \mathbf{v} \cdot \mathbf{x} = c\}$, where $\mathbf{v}$ is a fixed vector in $\mathbb{R}^d$ and $c \in \mathbb{R}$.

Consider the arrangement of 5 lines in Figure 1. How are they different? How are they the same? Is there a connection between hyperplane arrangements and polytopes?

An apparently unrelated topic...

Let G be an undirected graph. If we imagine G as a computer program, then the program gives us a directed graph called an *orientation* of G. In order to avoid infinite loops G should not have any directed circuits. How many *acyclic* orientations of G are there? What does this have to do with hyperplane arrangements? What does this have to do with graph coloring?

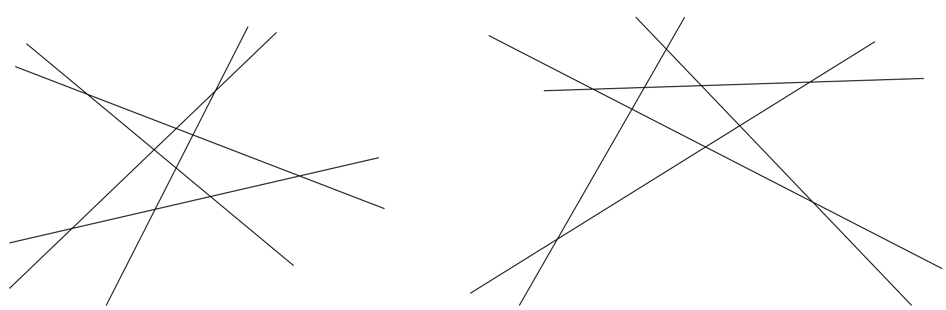

Figure 1. Two arrangements of 5 lines in general position

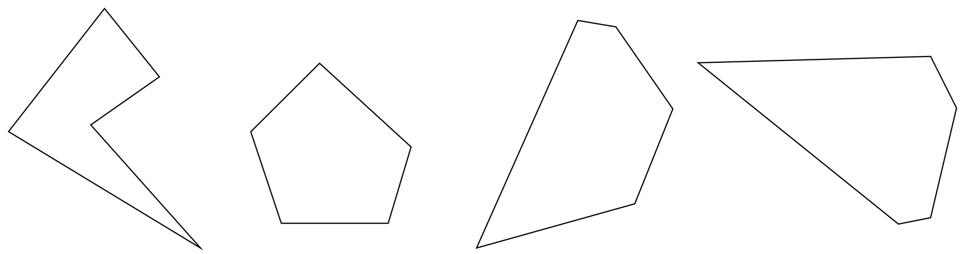

FIGURE 2. Four pentagons

The answers to the last several questions revolve around the Möbius function of a partially ordered set. This is an important and powerful tool for analyzing many combinatorial questions.

2. POLYTOPES

Recall that a convex d-polytope is a generalization of an n-gon. What exactly should that be? Before we answer that, perhaps we should examine what a polygon is!

In figure 2 we see four pentagons. All of these pentagons are *combinatorially* identical. What does that mean? The first pentagon is not convex. The others are all convex, but affinely inequivalent. What does that mean?

Definition 2.1. *A subset K of $\mathbb{R}^d$ is* **convex** *if for all $\mathbf{x}$ and $\mathbf{y}$ in K, the line segment from $\mathbf{x}$ to $\mathbf{y}$ is in K. Equivalently, for all $t \in [0,1], t\mathbf{x} + (1-t)\mathbf{y} \in K$.*

For a variety of reasons, we will only consider convex polygons and their generalizations. The *dimension* of a convex set is the dimension of its affine hull. There are several points of view one might take toward finding the analogue of a convex polygon. One is linear programming.

Example 2.2. *In Figure 3 we see the diamond as the feasible region in linear programming. Let x be the number of compact cars and y the number of SUV's that a factory makes in a given month. $x + y \geq a$, otherwise the factory is not making enough cars. $x + y \leq b$, since there are only so many workers in the plant. $x - y \leq c$ since the VP is really pushing SUV's. $x - y \geq d$ because of California emissions laws.*

Here is one possible generalization of polygons.

Definition 2.3. *P is a* **polyhedron** *if it is the intersection of a finite number of closed half-planes in $\mathbb{R}^d$. Equivalently, there exist $\mathbf{v}_1, \ldots \mathbf{v}_m \in \mathbb{R}^d$ and $c_1, \ldots, c_m \in \mathbb{R}$ such that $P = \{\mathbf{x} \in \mathbb{R}^d : \mathbf{v}_i \cdot \mathbf{x} \leq c_i$ for all $1 \leq i \leq m.\}$*

Even in dimension two this allows unbounded regions.

Definition 2.4. *An $\mathcal{H}$*-**polytope** *is a bounded polyhedron.*

Problem 2.5. *Show that the intersection of an $\mathcal{H}$-polytope with any hyperplane is an $\mathcal{H}$-polytope.*

Is this the right idea? Here is another possibility based on the notion of convex hull.

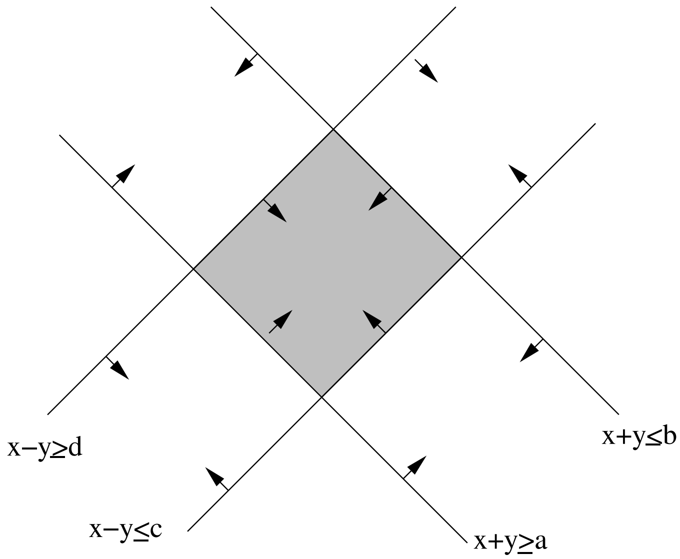

FIGURE 3. A diamond as a feasible region.

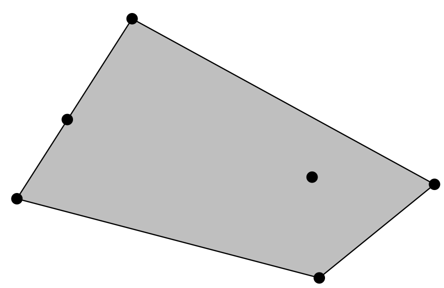

FIGURE 4. A $\mathcal{V}$-polytope

Definition 2.6. *Let $A \subseteq \mathbb{R}^d$. The smallest convex set containing A is called the* **convex hull** *of A.*

We use $ch(A)$ for the convex hull of A. Does this definition make sense? Yes, we can see that $ch(A)$ is the intersection of all convex sets K which contain Z.

Problem 2.7. *Let $A = \{\mathbf{x}_1, \ldots, \mathbf{x}_m\}$. Show that $ch(A) = \{\mathbf{y} : \mathbf{y} = \sum_{j=1}^{m} c_j \mathbf{x}_j, \sum_{j=1}^{m} c_j = 1,$ and every $c_j \geq 0.\}$ Show that in $\mathbb{R}^d$ one need only take convex sums of d points. What if A is not finite?*

It is easy to see that an n-gon is the convex hull of its vertices.

Definition 2.8. *P is a* **$\mathcal{V}$-polytope** *if it is the convex hull of a finite set of points.*

As can be seen in Figure 4, we do not insist that all of the points in the above definition be on the boundary.

Problem 2.9. *Let $T : \mathbb{R}^d \to R^{d'}$ be an affine map. Prove that if P is a $\mathcal{V}$-polytope, then $T(P)$ is also a $\mathcal{V}$-polytope. Prove that every $\mathcal{V}$-polytope is the affine image of the simplex.*

Which definition is correct? Is it obvious that problems 2.5 and 2.9 hold for the other definition?

Theorem 2.10. *P is an $\mathcal{H}$-polytope if and only if P is a $\mathcal{V}$-polytope.*

We postpone the proof of this fundamental result until Section 4.

Example 2.11. *The simplex, Δ^d is the convex hull of $e_1, \dots, e_{d+1}$ in R^{d+1}. The convex hull of any $d+1$ affinely independent points is combinatorially equivalent to Δ^d.*

The dimension of a polytope is the dimension of its affine hull. When P is d-dimensional we say P is a d-polytope. Every bounded closed d-dimensional convex set is homeomorphic to a d-ball, so the boundary of such a set is homeomorphic to a $(d-1)$-dimensional sphere. (What does this mean? How would you prove it?)

In order to discuss higher dimensional faces rigorously we introduce the important notion of a supporting hyperplane. Affine hyperplanes can be examined from several different points of view. The most obvious is as a point set, i.e., $H = \{(x_1, \dots, x_d) \in \mathbb{R}^d : c_1x_1 + \dots c_dx_d = a\}$, where $a, c_1, \dots, c_d$ are all scalars. We can also think of this H as the inverse image of a linear functional. Our previous H is $l_{\mathbf{v}}^{-1}(a)$, where $l_{\mathbf{v}}(\mathbf{x}) \equiv \mathbf{v} \cdot \mathbf{x}, \mathbf{v} = (c_1, \dots, c_d)$. In addition, H can be written as $\mathbf{y} + H'$, where $\mathbf{y} \in \mathbb{R}^d$ and H' is a linear hyperplane. Which, if any, of these descriptions are unique?

Definition 2.12. *Let $H = l_{\mathbf{v}}^{-1}(a)$ be an affine hyperplane. Then H is a supporting hyperplane of a polytope P if $l_{\mathbf{v}}(\mathbf{x}) \le a$ for all $\mathbf{x} \in P$ and $P \cap H \neq \emptyset$. A* **face** *of P is $\emptyset, P$ or any set of the form $P \cap H$, where H is a supporting hyperplane of P.*

This definition makes sense for any closed convex set K. For instance, all of the boundary points of a closed disk are faces of that disk.

Problem 2.13. *The d-cube is $\{(x_1, \dots, x_d) \in \mathbb{R}^d : -1 \le x_i \le 1.\}$ What are the faces of the d-cube? How many faces are there in each dimension? What are the faces of Δ^d?*

From problem 2.5 and Theorem 2.10 every face of a polytope is a polytope. The dimension of a face is its dimension as a polytope. The empty set (dimension $= -1$) and P are always faces of P. They are the improper faces of P. The highest dimensional proper faces of P are called the *facets* of P. Every face of P other than P is also a face of some facet and every face of a facet is also a face of P. The f-vector of a d-dimensional polytope P is $(f_0, f_1, \dots, f_{d-1})$, where f_i is the number of i-dimensional faces of P. Sometimes we think of $f_{-1} = f_d = 1$ and $f_i = 0$ for $i < -1$ and $i > d$. One of our main goals is to answer "What can you tell me about the f-vector of polytopes?"

3. Cyclic polytopes

Let $\mathbf{x}(t) = (t, t^2, \dots, t^d)$ be the moment curve in $\mathbb{R}^d$. Choose $n \ge d+1$ points $t_1 < t_2 < \dots < t_n$ in $\mathbb{R}$. Define $\mathbf{x}_i = \mathbf{x}(t_i)$ and let $C_d(n) \equiv ch(\{\mathbf{x}_1, \dots, \mathbf{x}_n\})$. For now, we will suppress the dependence of $C_d(n)$ on the t_i. We call $C_d(n)$ a *cyclic* polytope. Figure 5 shows a cyclic polytope in R^2.

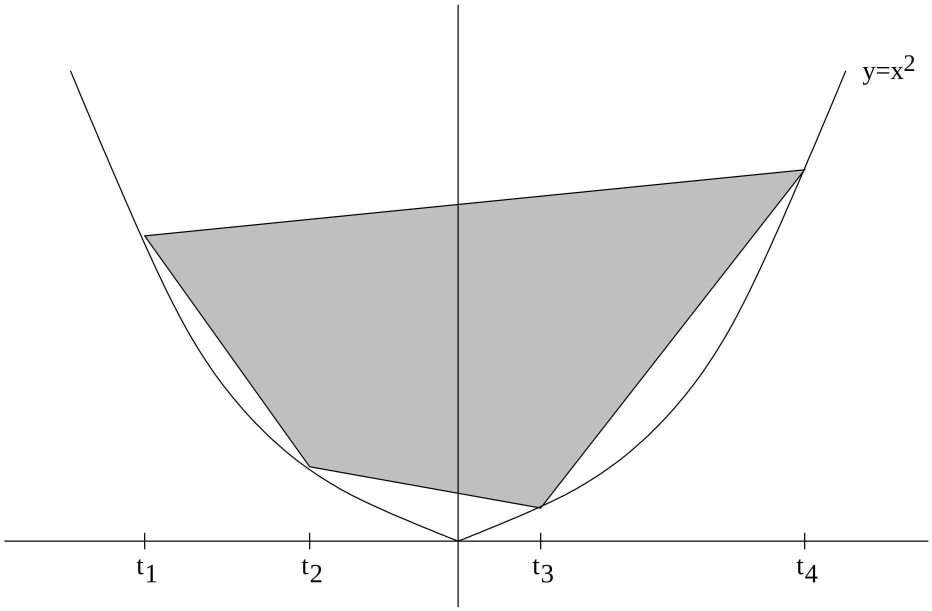

FIGURE 5. Cyclic polytope in R^2.

Proposition 3.1. *Every $d+1$-subset of $X = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ is affinely independent.*

Proof. - The proof is by induction on d, with $d = 1$ being obvious. Suppose the proposition holds for $d-1$. For notational simplicity we will show that $\mathbf{x}_1, \ldots, \mathbf{x}_{d+1}$ are affinely independent. The argument we give will clearly apply to any $(d+1)$-subset of X. Consider the matrix,

$$\begin{bmatrix} 1 & 1 & \ldots & 1 \\ t_1 & t_2 & \ldots & t_{d+1} \\ t_1^2 & t_2^2 & \ldots & t_{d+1}^2 \\ \vdots & \vdots & & \vdots \\ t_1^d & t_2^d & \ldots & t_{d+1}^d \end{bmatrix}. \tag{1}$$

It is sufficient to show that the rank of this matrix is $d+1$. For each $i > 1$, multiply row $i-1$ by t_1 and subtract it from row i. We obtain this matrix:

$$\begin{bmatrix} 1 & 1 & \ldots & 1 \\ 0 & (t_2 - t_1) \cdot 1 & \ldots & (t_{d+1} - t_1) \cdot 1 \\ 0 & (t_2 - t_1)t_2 & \ldots & (t_{d+1} - t_1)t_{d+1} \\ \vdots & \vdots & & \vdots \\ 0 & (t_2 - t_1)t_2^{d-1} & \ldots & (t_{d+1} - t_1)t_{d+1}^{d-1} \\ 0 & (t_2 - t_1)t_2^d & \ldots & (t_{d+1} - t_1)t_{d+1}^d \end{bmatrix}.$$

Now use the induction hypothesis to see that the lower right-hand $d \times d$ matrix has rank d and hence the entire matrix has rank $d+1$. $\square$

Problem 3.2. *Matrix (1) is known as the Vandermonde matrix. Compute its determinant.*

Corollary 3.3. *Every proper face of $C_d(n)$ is a simplex.*

Definition 3.4. *A polytope is* **simplicial** *if every proper face is a simplex.*

Example 3.5. *The d-crosspolytope is the convex hull of $\pm e_1, \ldots, \pm e_d$. It is a simplicial polytope. What are its faces? Can you compute f_k for this polytope?*

What are the facets of $C_d(n)$? Since this is a simplicial d-polytope we want to know which subsets of cardinality d of $X = \{\mathbf{x}_1, \ldots, \mathbf{x}_n\}$ are the vertices of a facet. (Why?) For any $S \subseteq [n]$ let X_S be the corresponding subset of X.

Theorem 3.6. [7] *If $S \subseteq [n]$ and $|S| = d$, then $ch(X_S)$ is a facet of P if and only if for every $1 \le i < j \le n$* **not** *in S,*

$$|k : i < k < j, k \in S|$$

is even.

Problem 3.7. *Using this, show that if $|S| \le d/2$, then $ch(X_S)$ is a face of P. Determine all of the faces of P.*

Notice that the combinatorial structure of P does not depend on the choice of the t_i. We also note that the above problem says that the f-vector for P is as large as possible for $i \le d/2$. What about the number of higher dimensional faces of $C_d(n)$? Amazingly, we will see that for *any* simplicial d-polytope, once we know $f_0, \ldots, f_{\lfloor d/2 \rfloor}$, we know all of the f_i.

Proof. (Gale) Let $X_S = \{\mathbf{x}_{i_1}, \ldots, x_{i_d}\}$. Let $l_S(\mathbf{x})$ be defined by

$$l_S = \det \begin{bmatrix} 1 & 1 & 1 & \ldots & 1 \\ \mathbf{x} & \mathbf{x}(t_{i_1}) & \mathbf{x}(t_{i_2}) & \ldots & \mathbf{x}(t_{i_d}) \end{bmatrix}.$$

Using the permutation formula for determinants, we see that $l_S = l_{\mathbf{v}} - c$, where $l_{\mathbf{v}}$ is a linear functional and $c \in \mathbb{R}$. Furthermore, $l_S = 0$ on X_S. (Why?) So, the only possible supporting hyperplane for a simplex with vertices X_S is the affine hyperplane $l_S = 0$. When do all the points of P lie on one side of l_S? Consider the following polynomial.

$$f_S(t) = \det \begin{bmatrix} 1 & 1 & 1 & \ldots & 1 \\ \mathbf{x}(t) & \mathbf{x}(t_{i_1}) & \mathbf{x}(t_{i_2}) & \ldots & \mathbf{x}(t_{i_d}) \end{bmatrix}.$$

Now, $f_S(t)$ is a polynomial of degree at most d with d distinct zeros $t_{i_1}, \ldots, t_{i_d}$. Therefore, $f_S(t)$ changes sign every time it passes through one of the t_{i_j}. So, in order for $l_S = 0$ to be a supporting hyperplane it is necessary and sufficient that the number of points in S be even between every pair of points in $X - X_S$. □

Problem 3.8. *Verify that $C_3(6)$ and the octahedron have the same f-vector. Are they combinatorially equivalent?*

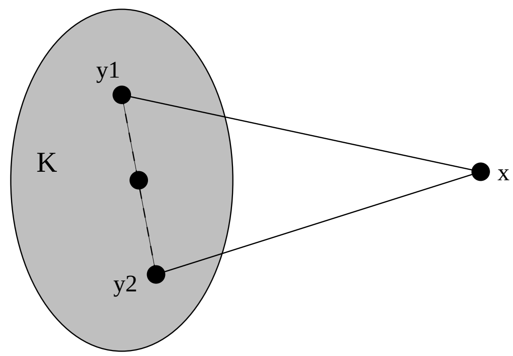

FIGURE 6. Uniqueness of projection of x to K.

4. Proof of main theorem

First, several preliminary results. We use $\delta(\mathbf{x}, \mathbf{y})$ for the distance between two points. Our proof closely follows the relevant sections of “Convex Polytopes”, B. Grünbaum.

Lemma 4.1. *Let K be a closed convex set and suppose $\mathbf{x} \notin K$. Then there exists a unique point $\mathbf{y} \in K$ which minimizes $\delta(\mathbf{x}, \mathbf{y})$ over all $\mathbf{y} \in K$.*

Proof. Since K is closed there is at least one $\mathbf{y}_1$ with the stated property. Suppose $\mathbf{y}_2$ is another point of K which minimizes the distance to $\mathbf{x}$. Then, convexity of K implies that the midpoint of the line segment from $\mathbf{y}_1$ to $\mathbf{y}_2$ is in K, while elementary geometry shows that this point would be closer to $\mathbf{x}$. □

Proposition 4.2. *Let K and K' be closed convex sets in $\mathbb{R}^d$ with K bounded. If $K \cap K' = \emptyset$, then there exists a hyperplane which separates K and K'.*

Proof. For each $\mathbf{y} \in K$ let $\delta_{K'}(\mathbf{y})$ be the distance from $\mathbf{y}$ to K'. Since $\delta_{K'}$ is continuous and K is compact, there is a $\mathbf{y}_0 \in K$ which minimizes $\delta_{K'}(\mathbf{y})$ on K. By the above lemma, there is a *unique* point $\mathbf{x}_0 \in K'$ such that $\delta_{K'}(\mathbf{y}_0) = \delta(\mathbf{x}_0, \mathbf{y}_0)$. Let $H_{\mathbf{x}}$ and $H_{\mathbf{y}}$ be the hyperplanes perpendicular to the line segment $[\mathbf{x}_0, \mathbf{y}_0]$ which intersect at $\mathbf{x}_0$ and $\mathbf{y}_0$ respectively. Claim: Any hyperplane H parallel to $H_{\mathbf{x}}$ and $H_{\mathbf{y}}$ which intersects the open line segment works. This is clear as long as there are no points of K or K' in the open slab between $H_{\mathbf{x}}$ and $H_{\mathbf{y}}$. Let $\mathbf{z}$ be such a point. As can be seen in Figure 7, if $\mathbf{z} \in K$, then points near $\mathbf{y}_0$ on the line segment from $\mathbf{y}_0$ to $\mathbf{z}$ will be closer to K' than $\mathbf{y}_0$. Similarly, if $\mathbf{z} \in K'$ then points near $\mathbf{x}_0$ on the line segment from $\mathbf{x}_0$ to $\mathbf{z}$ will be closer to $\mathbf{y}_0$ than $\mathbf{x}_0$. □

Problem 4.3. *What if K is unbounded?*

Problem 4.4. *Prove that every closed convex set is the intersection of all the closed half-spaces which contain it.*

Proposition 4.5. *Let $\mathbf{x}$ be in the boundary of a d-dimensional closed convex set $K \subseteq \mathbb{R}^d$. Then $\mathbf{x}$ is contained in a face of K.*

Proof. W.L.O.G. we assume that $\mathbf{0}$ is in the interior of K. Let $K_\epsilon = \{\epsilon\mathbf{x} : \mathbf{x} \in K\}$. Since $\mathbf{x} \in \partial K, \mathbf{x} \notin K_{(1-\epsilon)}$ for $0 < \epsilon < 1$. (Why? - This is a point fudged by

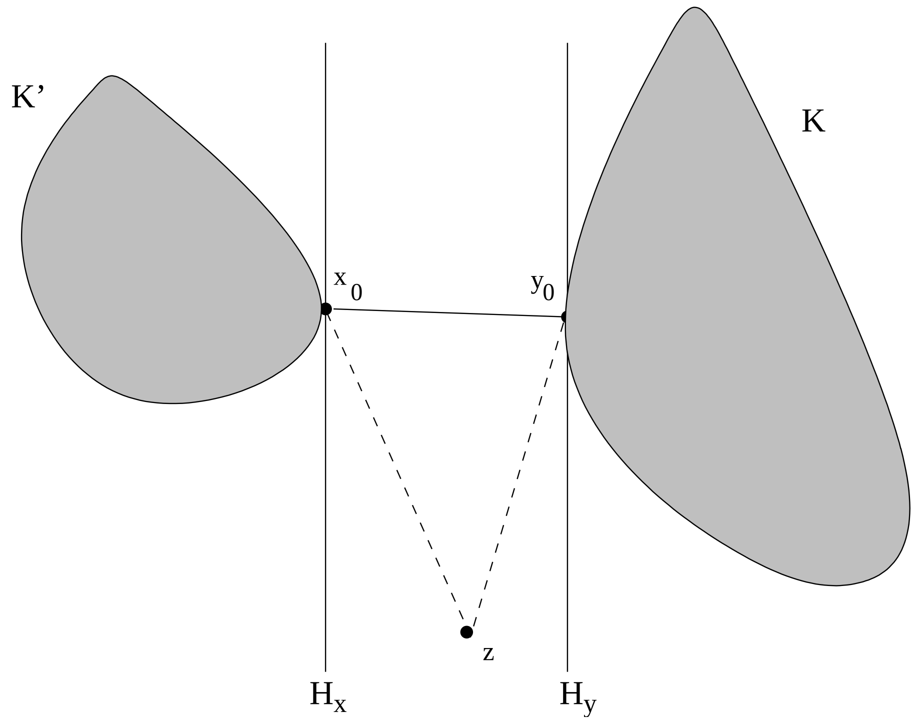

FIGURE 7

just about everyone...) By the previous proposition, for each such ϵ there exist a hyperplane H_ϵ which separate $\mathbf{x}$ and $K_{(1-\epsilon)}$. The colleciton of all the H_ϵ must have at least one limit plane H which will satisfy the proposition. (What is a limit hyperplane? Why does one exist? Why does it work?) □

Just as in the two- and three-dimensional cases, we will call a zero-dimensional face of K a *vertex*. Not surprisingly, the vertices of P play a critical role in the proof of Theorem 2.10. However, it turns out that it is easier to work with the closely related notion of extreme points.

Definition 4.6. $\mathbf{x}$ *is an extreme point of a convex set* K *if* $\mathbf{x}$ *is never in the relative interior of a line segment contained in* K. *We denote the extreme points of* K *by* $\operatorname{ext} K$.

Lemma 4.7. *Let* K *be convex.*

- *Every vertex of* K *is an extreme point of* K.
- *If* $K = ch(A)$, *then* $\operatorname{ext} K \subseteq A$.
- *If* F *is a face of* K, *then* $\operatorname{ext} F = F \cap \operatorname{ext} K$.

Problem 4.8. *Prove the lemma.*

While vertices and extreme points are the same in polyhedra, this is not true for general convex sets. For example, in Figure 8, F_2 is a face of F_1 and F_1 is a face of K, but F_2 is not a face of K.

Theorem 4.9. *If* K *is a compact convex set, then* $ch(\operatorname{ext} K) = K$.

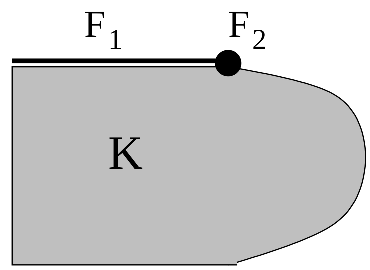

FIGURE 8. Faces of faces may not be faces!

Proof. Evidently, $ch(\operatorname{ext} K) \subseteq K$. We prove the reverse inclusion by induction on the dimension of K. Dimension one is obvious.

Suppose that $\mathbf{x} \in K$. If $\mathbf{x}$ is an extreme point then obviously $\mathbf{x} \in ch(\operatorname{ext} K)$. So, let $[\mathbf{y}_0, \mathbf{z}_0]$ be a line segment in K containing $\mathbf{x}$ in its relative interior. Extend this segment in both directions. It will intersect the boundary of K at two points, $\mathbf{y}$ and $\mathbf{z}$. By the previous proposition there exist faces $F_{\mathbf{y}}$ and $F_{\mathbf{z}}$ which contain $\mathbf{y}$ and $\mathbf{z}$ respectively. Each of these faces has dimension less that K, so by the induction hypothesis, $F_{\mathbf{y}} = ch(\operatorname{ext} F_{\mathbf{y}})$ and $F_{\mathbf{z}} = ch(\operatorname{ext} F_{\mathbf{z}})$. As $\mathbf{x} \in ch(\operatorname{ext} F_{\mathbf{y}} \cup \operatorname{ext} F_{\mathbf{z}})$ and the lemma tells us that $\operatorname{ext} F_{\mathbf{y}} \cup \operatorname{ext} F_{\mathbf{z}} \subseteq \operatorname{ext} K$, we are done.

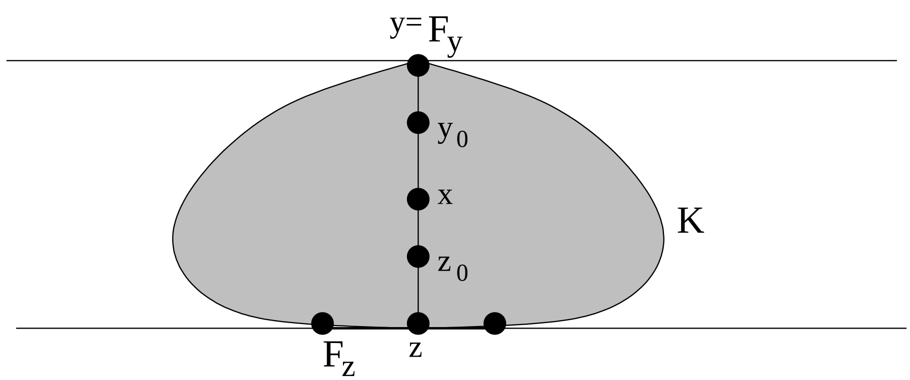

FIGURE 9. $ch(\operatorname{ext} K) = K$

□

We are finally ready to prove Theorem 2.10.

Let P be an $\mathcal{H}$-polytope. We prove P is a $\mathcal{V}$-polytope by induction on dimension. As usual, dimension one is trivial.

Write P as a minimal intersection of closed half-spaces $P = \cap_{i=1}^m H_i$. By minimal we mean that if we set $P_i = \cap_{j \neq i} H_j$, then $P_i \neq P$ for any i. Minimality guarantees that the boundary of P is contained in the union of the $(d-1)$-dimensional faces of P and there are m of these. (Why?) By induction, each such face has a finite number of extreme points. All of the extreme points of P are on the boundary. Hence, the previous lemma and the last theorem imply that P has a finite number of extreme points and must be a $\mathcal{V}$-polytope.

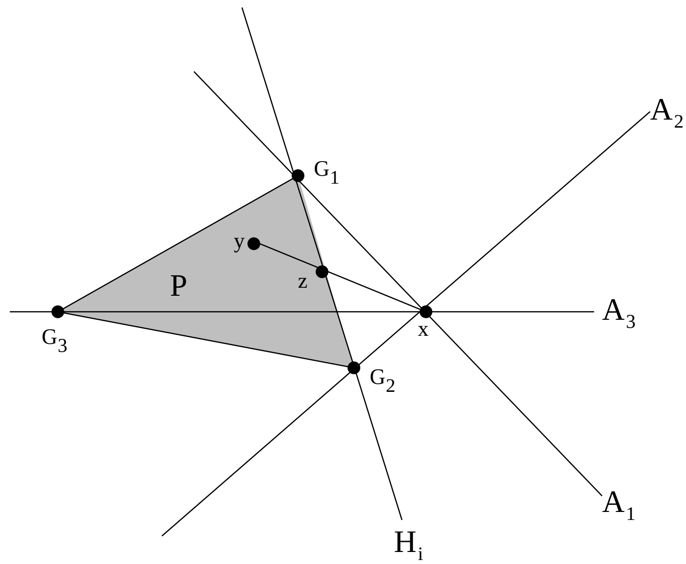

FIGURE 10. $P = \cap H_i$

Now suppose that P is a d-dimensional $\mathcal{V}$-polytope in $\mathbb{R}^d$. Write $P = ch(V), |V| = n$. By the previous theorem, we might as well assume that $V = \text{ext}\, P$. Any k-face of P is determined by $(k+1)$ affinely independent points in the face. Hence, for every k, $f_k \leq \binom{|V|}{k+1}$. Let F_i enumerate the facets (the $(d-1)$-dimensional faces) of P and let H_i be the closed half-spaces of their corresponding supporting hyperplanes. The proof is complete once we show the following.

Claim: $P = \cap H_i$.

Proof. (of claim) Certainly $P \subseteq \cap H_i$. Suppose $\mathbf{x} \notin P$. For each face G_j of P of dimension $d-2$ or less, let A_j be affine hull of G_j and $\mathbf{x}$. The A_j form a finite collection of affine subspaces each of which has dimension at most $d-1$. Therefore their union does not cover the interior of P. So, there exists $\mathbf{y}$ in the interior of P so that $\mathbf{y}$ is not in the union of the A_j. Consider the line segment from $\mathbf{y}$ to $\mathbf{x}$. Since $\mathbf{y} \in P$ and $\mathbf{x} \notin P$ there exists $\mathbf{z}$ on the boundary of P and the line segment from $\mathbf{y}$ to $\mathbf{x}$. By Proposition 4.5 there is a supporting hyperplane of P which contains $\mathbf{z}$. As $\mathbf{y}$ is not in any of the A_j, it must be the case that $\mathbf{z}$ is in one of the H_i. Hence, $\mathbf{x} \notin \cap H_i$. □

5. SHELLING

The boundary of a d-polytope is homeomorphic to the $(d-1)$-sphere. It turns out that this boundary can be built up by gluing the facets together in a "nice" way. This will lead us to some remarkable enumerative results for the f-vector of the polytope, especially when P is simplicial. Until further notice, P is a simplicial d-polytope in R^d.

A *shelling* of P is an ordering $F_1, \ldots, F_m$ of the facets of P such that for all $j \geq 2, F_j \cap (\cup_{i=1}^{j-1} F_i)$ is a nonempty union of facets of F_j. If P has a shelling order, then we say P is *shellable.* Figure 11 shows a shelling of the boundary of the

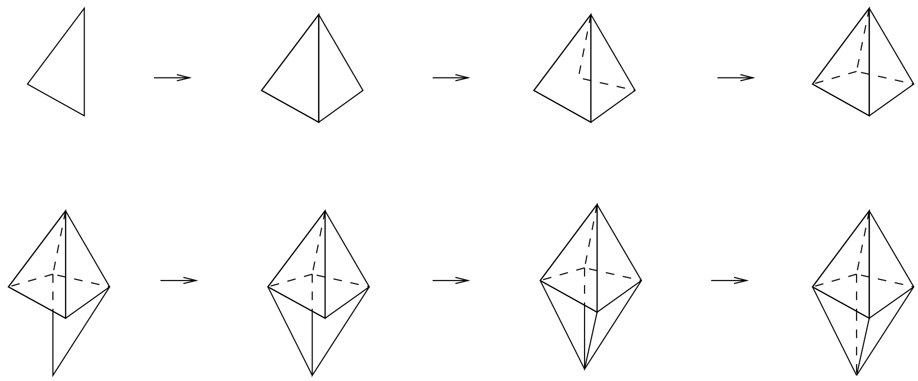

FIGURE 11. A shelling of the boundary of the octohedron

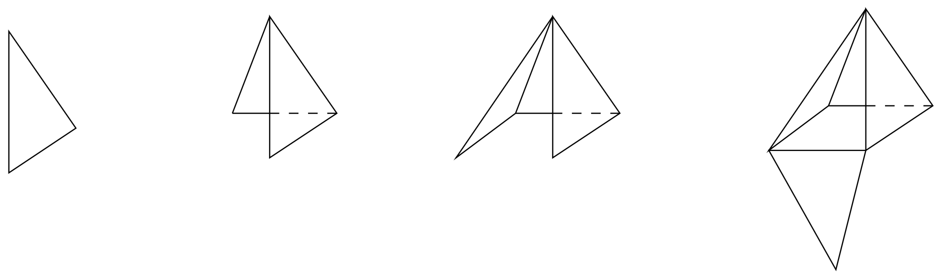

FIGURE 12. The beginning of a non-shelling of the boundary of the octohedron

octrahedron, while figure 12 shows an ordering of the four facets of the boundary of the octohedron which is not the beginning of any shelling.

Now view the ordering of the facets as a way of building up ∂P. Each facet adds new faces of ∂P as we take the union of the facet with the previous ones. The shelling condition insures that there is a unique new minimal face added at each step. Indeed, at the j^{th} step $F_j \cap (\cup_{i=1}^{j-1} F_i)$ is a union of facets of F_j. The minimal face is $M_j = ch(\{v_{j_1}, \ldots, v_{j_m}\})$, where the v_{j_k} are the vertices opposite the facets of F_j in the intersection. Figure 13 shows the minimal new faces of the shelling in Figure 11.

The concept of shellability is very important and extends to a number of situations including abstract simplicial complexes. An *abstract simplicial complex* consists of a set V, the vertices of the complex, and a set of faces $\Delta \subseteq 2^V$ The faces must be closed under subsets. If $F \in \Delta$ and $G \subseteq F$, then $G \in \Delta$. The maximal faces of Δ are the *facets* of the complex. The dimension of a simplex $G \in \Delta$ is $|G| - 1$.

Example 5.1. *Some examples of abstract simplicial complexes.*

(1) *Boundaries of simplicial polytopes.*
(2) *Simple graphs.*
(3) *Let V be a subset of vectors in a vector space and let Δ be the subsets of V which form independent subsets of vectors. What are the facets of this complex?*

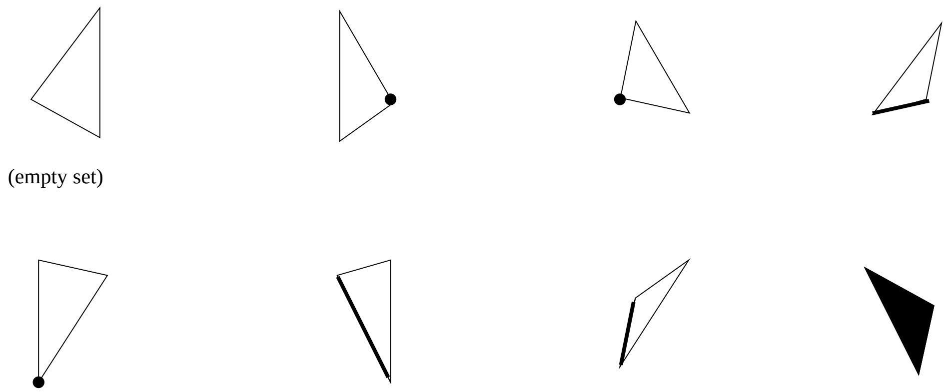

FIGURE 13. Minimal new faces of the shelling

Every abstract simplicial complex has a geometric realization $|\Delta|$. Let $[n] = |V|$. Then $|\Delta|$ is the subset of $\mathbb{R}^n$ defined by

$$|\Delta| = \cup_{F \in \Delta} ch(\{e_i\}_{i \in F}).$$

The definition of shellable for an abstract simplicial complex is exactly the same as for simplicial polytopes; for $j \geq 2, F_j \cap (\cup_{i=1}^{j-1} F_i)$ must be a union of facets of the boundary of F_j.

Problem 5.2. *Show that connected graphs and the complexes described in (3) are shellable.*

Example 5.3. *Let G be the graph in Figure 14. Let the vertices of Δ be the edges of the graph. The faces of Δ are those subsets of edges whose removal does not disconnect the graph. Specifically, $\emptyset$, all singleton, and all doubletons except $\{a, b\}$ and $\{c, d\}$ are the faces of Δ. We can project $|\Delta|$ down to $\mathbb{R}^2$ to get the complex in Figure 14. Ordering the facets as shown, we see that $\emptyset$ is the minimal new face for* 1, *and the minimal new face is a single vertex for faces* 2, 3, 4 *and an edge for faces* 5, 6, 7, 8.

Suppose that G above represents a network and each edge has equal and independent probability of failing $p, 0 < p < 1$. What is the probability that that network will remain connected? Directly checking each possible subset of edges which keep the graph connected we see that this probability is

$$(1-p)^5 + 5p(1-p)^4 + 8p^2(1-p)^3 = (1-p)^3[1 + 3p + 4p^2].$$

Hmmm....the coefficient in each degree i of the last factor is the number of steps in the shelling in which we added a minimal new face of cardinality i.

For each i, let h_i be the number of facets whose unique new minimal face has cardinality i. The *h-vector* of ∂P is $(h_0, \dots, h_d)$.

Example 5.4. *The shelling in Figure 11 gives h-vector* $(1, 3, 3, 1)$.

The *shelling polynomial* of the shelling $F_1, \dots, F_m$ is

$$h_{\partial P}(x) = h_0 x^d + h_1 x^{d-1} + \cdots + h_{d-1} x + h_d. \tag{2}$$

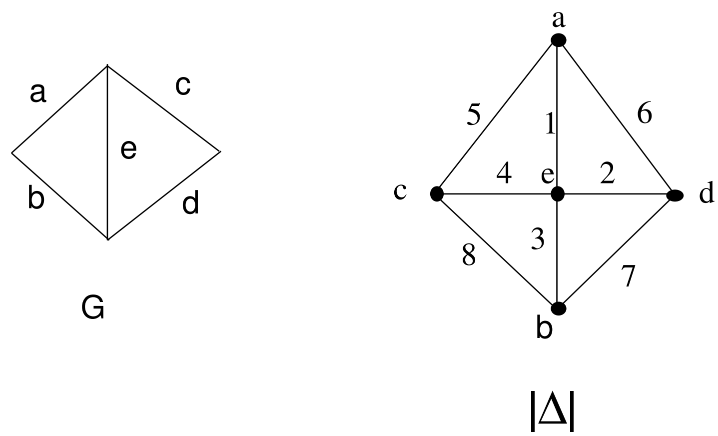

FIGURE 14. G and $|\Delta|$

The shelling polynomial appears to depend on the shelling. We can also encode the f-vector of P in a polynomial. Define the *face polynomial* of ∂P to be

$$f_{\partial P}(x) = f_{-1}x^d + f_0 x^{d-1} + f_1 x^{d-2} + \cdots + f_{d-1}.$$

Theorem 5.5. $h_{\partial P}(x+1) = f_{\partial P}(x)$.

Proof. Suppose that at a particular shelling step the minimal new face has cardinality i. Then this step adds $\binom{d-i}{0}$ to f_i, $\binom{d-i}{1}$ to f_{i+1} and in general $\binom{d-i}{j}$ to f_{i+j} up to $j = d-i$. This has the same effect as adding $(x+1)^{d-i}$ to $f_{\partial P}(x)$. Summing up over all the shelling steps gives the required formula. $\square$

Corollary 5.6. *The h-vector of ∂P does not depend on the shelling order.*

Problem 5.7. *Write down formulas for h_i in terms of f_j and vice versa. Which f_j does h_i depend on (and vice versa)? Show that if $\mathcal{P}$ is a collection of simplicial d-polytopes and $P \in \mathcal{P}$ has the property that $h_i(\partial P) \geq h_i(\partial P')$ for all $0 \leq i \leq d$ and $P' \in \mathcal{P}$, then $f_i(\partial P)$ also maximizes all $f_i(\partial P')$ in $\mathcal{P}$. Is this still true if we reverse the role of the f- and h-vector?*

Even if ∂P does not have a shelling we can still define $h_{\partial P}(x) = f_{\partial P}(x-1)$ and define the h-vector according to (2). For instance, for any simplicial 3-polytope we know that $f_2 = 2/3 f_1$ and $f_0 - f_1 + f_2 = 2$, so $f_0 - f_1 + 2/3 f_1 = 2 \rightarrow f_1 = 3 \cdot f_0 - 6$ and $f_2 = 2 \cdot f_0 - 4$. This implies that $h_0 = 1, h_1 = f_0 - 3 = h_2$ and $h_3 = 1$.

Does every polytope have a shelling? In fact something even better is true. Every polytope has a *line* shelling. These were introduced by Bruggesser and Mani [3]. This shelling can be described as follows. Let $\mathbf{x}$ be a point in the interior of P. Now choose a line r through $\mathbf{x}$ with the following two properties:

(1) r intersects every supporting hyperplane of the facets of P.
(2) r does not intersect any *non-trivial* intersection of the supporting hyperplanes of the facets of P.

Why is it obvious we can always choose such a line? Now imagine you are in a rocket ship inside the planet P heading along the line r in a chosen positive

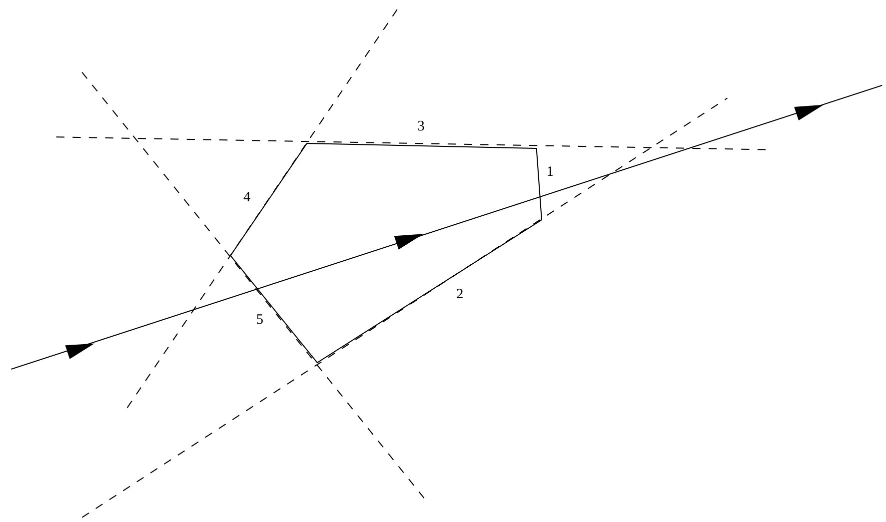

Figure 15. Line shelling of a polygon

direction. Initially you can not see anything as you are inside the planet. By (2) you will emerge at a facet, F_1. As you travel away from P you will see one new facet each time you pass through its corresponding supporting hyperplane. The shelling order for these facets is the same as the order you see them. By (2) you will never see two (or more) new facets simultaneously. Eventually you will go far enough towards "$+\infty$" so that you can see as many facets as possible in that direction. Now begin your return to P coming from the "$-\infty$" direction. At this point you can see all the facets you could not see from the "$+\infty$" direction. As you pass through each of their supporting hyperplanes the corresponding facet will disappear from your vision. The shelling order continues from before in the order in which the facets disappear from your vision. By (1) each facet of P occurs exactly once in your shelling order.

Problem 5.8. *Show that this is a shelling of P. Where did you use convexity?*

The fact that every simplicial polytope is shellable already shows there are strong restrictions on their f-vectors. For instance, $(6, 15, 18, 7)$ is not the f-vector of any simplicial 3-polytope. Indeed, this would give $(1, 2, 3, 2, -1)$ as its h-vector, and this is impossible. In fact, for the same reason, $(6, 15, 18, 7)$ is not the f-vector of any shellable abstract simplicial complex. (What is the f-vector of an abstract simplicial complex?)

What happens if we travel along r in the opposite direction? This simply reverses the shelling order. But now, each facet which originally contributed to h_i, contributes to h_{d-i}. Since the h-vector is independent of the shelling order we obtain the *Dehn-Sommerville* equations.

Theorem 5.9. [6], [15] *If P is a simplicial d-polytope, then $h_i = h_{d-i}$.*

This is an even stronger restriction on the h-vector of P. Combining this theorem with Problem 5.7, we now see why once we know $f_0, \dots, f_{\lfloor d/2 \rfloor}$, we know the entire f-vector of P.

Are there any other restrictions on the h-vectors of simplicial polytopes? In order to discuss this we must introduce some notation.

Let a, i be natural numbers. Then there is a unique way of writing

$$a = \binom{a_i}{i} + \binom{a_{i-1}}{i-1} + \cdots + \binom{a_j}{j}, a_i > a_{i-1} > \cdots > a_j \geq j.$$

Problem 5.10. *Prove the above statement.*

Example 5.11. $a = 14, i = 3$. *Then* $14 = \binom{5}{3} + \binom{3}{2} + \binom{1}{1}$.

With a decomposed as above define

$$a^{<i>} = \binom{a_i + 1}{i+1} + \binom{a_{i-1} + 1}{i} + \cdots + \binom{a_j + 1}{j+1}.$$

Example 5.12. *So,* $14^{<3>} = \binom{6}{4} + \binom{4}{3} + \binom{2}{2} = 15 + 4 + 2 = 21$.

In 1971 P. McMullen conjectured that necessary and sufficient conditions for $(h_0, \dots, h_d)$ to be the h-vector of the boundary of a simplicial d-polytope are

- $h_i \geq 0$.
- $h_i = h_{d-i}$.
- $h_0 \leq h_1 \leq \cdots \leq h_{\lfloor d/2 \rfloor}$.
- For $i \leq d/2$, define $g_i = h_i - h_{i-1}$. Then for $i \leq d/2$, $g_{i+1} \leq g_i^{<i>}$.

Stanley [18] proved necessity, while Billera and Lee [1] proved sufficiency. These results give a complete characterization of h-vectors (and hence f-vectors) of the boundary of simplicial polytopes. For instance, here are two problems that are now fairly easy.

Problem 5.13. *Show that $C_d(n)$ maximizes every f_i among all possible simplicial d-polytopes with n vertices. What is fewest number faces in dimension i that a simplicial d-polytope with n vertices can have?*

Looking at many examples, it appears that the f-vector of a simplicial polytope is unimodal. A sequence $(f_0, f_1, \dots, f_d)$ is *unimodal* if there exists i such that $f_0 \leq f_1 \leq \cdots \leq f_i \geq f_{i+1} \geq \cdots \geq f_d$). By Stanley's proof of the necessity of McMullen's conditions, the h-vector of a simplicial polytope is symmetric and unimodal. It can be shown that the f-vector of a simplicial d-polytope is unimodal when $d \leq 19$. Shortly after Billera and Lee proved the sufficiency of McMullen's conditions, Björner [2] and Lee [13] gave examples of 20-dimensional simplicial polytopes with trillions of vertices with $f_{11} > f_{12} < f_{13}$.

What about non-simplicial polytopes?

First we extend the definition of shelling to arbitrary polytopes as follows. An ordering $F_1, \dots, F_m$ of the facets of P (an arbitrary polytope) is a shelling of ∂P if for each $j \geq 2, F_j \cap (\cup_{i=1}^{j-1} F_i)$ is a union of facets of F_j which is an initial segment of a shelling of ∂F_j. Figure 16 shows the beginning of a shelling of the boundary of the cube.

The line shelling from before also gives a shelling of an arbitrary polytope with properties similar to the simplicial case. These new shellings allow us to prove Euler's formula for all polytopes.

Problem 5.14. *Prove that a line shelling is a shelling of an arbitrary polytope.*

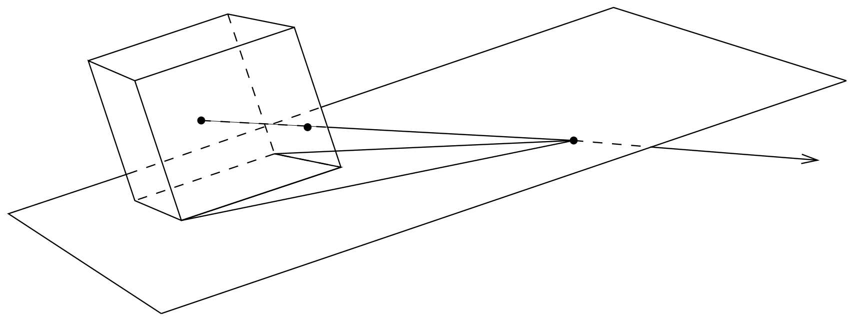

FIGURE 16. Shelling the boundary of a cube.

Definition 5.15. *The Euler characteristic of P is* $\chi(P) = \displaystyle\sum_{j=0}^{d} (-1)^j f_j$.

Consider the situation when P is simplicial. From the Dehn-Sommerville equations we know that $h_d = 1$. Therefore,

$$1 = h_{\partial P}(0) = f_{\partial P}(-1) = (-1)^d + (-1)^{d-1}[\chi(P) - (-1)^d].$$

This easily reduces to $\chi(P) = 1$.

Problem 5.16. *Show that $\chi(P) = 1$ for all polytopes. (Hint: Induction on dimension and on the line shelling. Along the way, show that until the last facet is included, partial shellings of ∂P also have Euler characteristic one.)*

There is an analogue of the Dehn-Sommerville equations for arbitrary polytopes. We will not discuss it here, other than to say that it involves more than just the f-vector of the polytope. It also includes the combinatorics of the faces of P as expressed in the face lattice of P.

The search to understand f-vectors of various complexes is an ongoing and extensive area of research. We mention just two results of Stanley which arise naturally in our setting. The Dehn-Sommervile equations hold for *any* simplicial complex whose geometric realization is homeomorphic to a sphere. Also, $h_{i+1} \leq h_i^{<i>}$ for any shellable simplicial complex. On the other hand, a major open question is whether or not McMullen's g-conditions hold for h-vectors of arbitrary simplicial complexes whose geometric realizations are homeomorphic to a sphere.

6. The face poset

One of the fundamental concepts in combinatorics (mathematics?) is a partially ordered set.

Definition 6.1. *A poset, or partially ordered set, is a set L with a binary relation $\leq$ which is anti-symmetric, transitive and reflexive.*

As usual, we use $x < y$ to indicate that $x \leq y$ and $x \neq y$.

Example 6.2. *Some examples*

- *Usual $\leq$ on $\mathbb{R}$ or $\mathbb{Z}$.*
- *$a \leq b$ if and only if $a|b$ in $\mathbb{Z}^+$.*

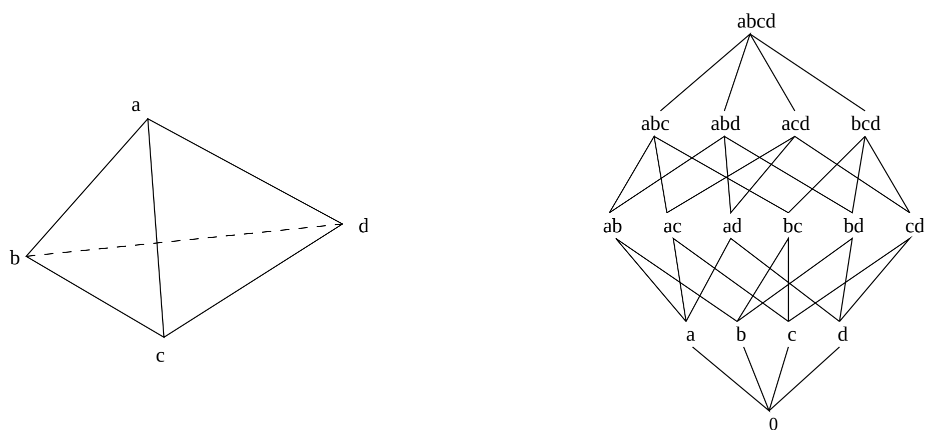

FIGURE 17. Hasse diagram of $\mathcal{F}(\Delta^3)$

- *$\mathcal{F}$ any collection of subsets of any fixed set S. Define $A \le B$ iff $A \subseteq B$. When $\mathcal{F}$ is the faces of a polytope P, we call it the face poset of P and denote it $\mathcal{F}(P)$.*

When the binary relation $\le$ is unambiguous we will suppress it and simply write L is a poset.

Two posets $(L, \le_L)$ and $(M, \le_M)$ are *isomorphic* if there is a bijection $\phi : L \to M$ such that $x \le_L y$ if and only if $\phi(x) \le_M \phi(y)$. In this case ϕ is called a poset isomorphism. Two polytopes P and Q are *combinatorially equivalent* if $\mathcal{F}(P)$ is isomorphic to $\mathcal{F}(Q)$.

Finite posets can be represented by their Hasse diagram. An edge between two elements x and y of the poset in the Hasse diagram indicates that the upper element (say x) *covers* the lower element. That is, $y \le x$ and there is no z in the poset such that $z \ne x, z \ne y$ and $y \le z \le x$. Figure 17 shows the Hasse diagram of $\mathcal{F}(\Delta^3)$.

What can we say about $\mathcal{F}(P)$?

Proposition 6.3. *If $F \in \mathcal{F}(P)$ and $G \in \mathcal{F}(F)$, then $G \in \mathcal{F}(P)$.*

Proof. Let H_F be a supporting hyperplane of F. Let H_G be a supporting hyperplane of G in H_F. Now rotate H_F along H_G in a direction away from the vertices of F not in G as in Figure 18. □

Thus, the *lower interval* $[\emptyset, F] = \{G \in \mathcal{F}(P) : \emptyset \subseteq G \subseteq F\}$ equals $\mathcal{F}(F)$. What do you think upper intervals, $[F, P] = \{G :\in \mathcal{F}(P), F \subseteq G \subseteq P\}$, look like?

Here are some other properties of $\mathcal{F}(P)$ you should try to prove.

Proposition 6.4. *Let P be a d-polytope.*

(1) *If F is a $(d-2)$-face of P, then F is the intersection of two facets of P.*
(2) *If F is a $(d-k)$-face of P, then F is the intersection of k facets of P.*
(3) *The intersection of two faces of P is a face of P.*
(4) *If F is a face of P other than P, then F is a face of some facet of P.*

A *chain* in a poset L is a sequence $c = x_0 < x_1 < \cdots < x_r$ in L. The *length* of c is r. A poset is *ranked* (also called *graded*) if every maximal chain in L has the same length. If L is ranked, then the rank of $x \in L$ is the length of the longest chain in

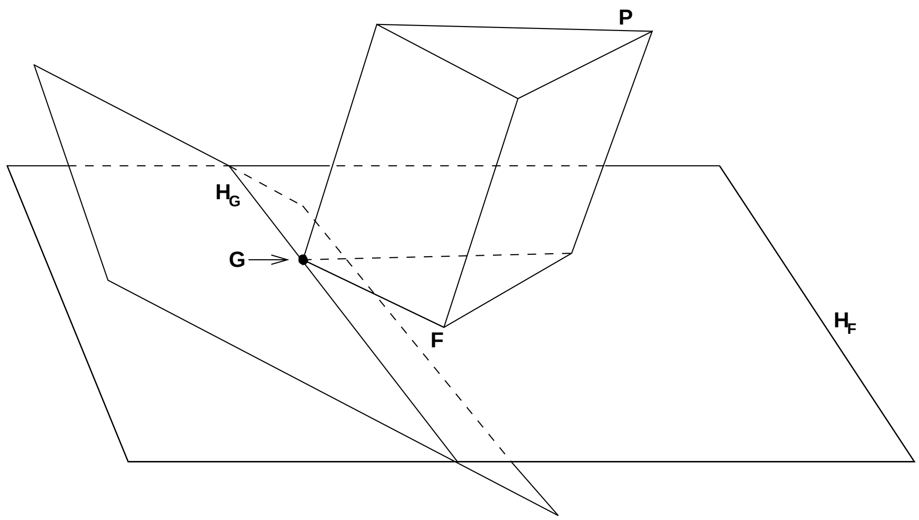

FIGURE 18

L in which x is the last element. From the above Proposition we see that $\mathcal{F}(P)$ is ranked and that the rank of a face in $\mathcal{F}(P)$ is its dimension plus one.

Let $x, y \in L$. Their *meet* is their greatest lower bound and is denoted $x \wedge y$. Similarly, their *join* is their least upper bound and is denoted $x \vee y$. Of course, there is no reason that L should contain a meet or join for x and y. We can see that the face poset of a polytope contains the meet of any two faces, namely their intersection. What about the join?

Proposition 6.5. *Let L be a finite poset with a greatest element. Suppose that $x \wedge y$ exists in L for every pair of elements x, y in L. Then, $x \vee y$ exists for every pair of elements in L.*

Proof. By induction every finite subset of L has a greatest lower bound. Given $x, y \in L$, let $U = \{z \in L : x \leq z, y \leq z.\}$. Since L has a greatest element, U is finite and not empty. Now check that the greatest lower bound of U is $x \vee y$. $\square$

Any poset which contains $x \vee y$ and $x \wedge y$ for all x and y in the poset is called a *lattice*. For this reason $\mathcal{F}(P)$ is also called the face lattice of P.

Problem 6.6.

(1) *Find an infinite poset with a greatest element such that $x \wedge y$ exists for every x, y, but $x \vee y$ does not exist.*

(2) *Which of the following posets are lattices? Which are ranked?*

 (a) *Subspaces of a fixed vector space with $\subseteq$.*

 (b) *Subsets of an $m \times n$ chess board which are the squares of consecutive knight moves which begin in the lower left-hand corner and do not repeat any square. Include the lower left-hand corner as the square visited by a sequence of zero knight moves.*

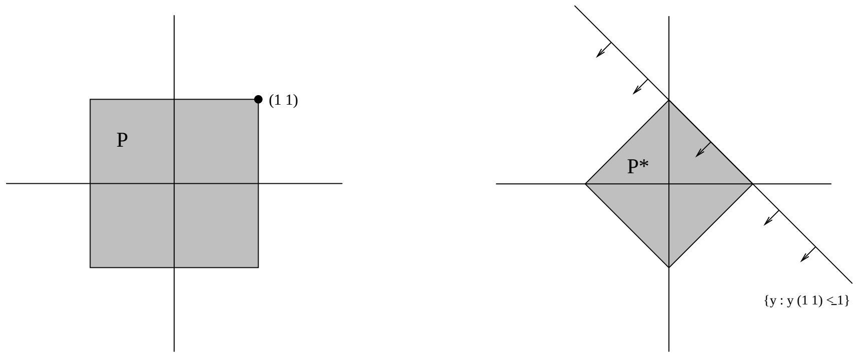

FIGURE 19. Polar of the square

(3) *Let P be a d-polytope in $\mathbb{R}^d$. Embed P in $\mathbb{R}^{d+1}$ by embedding all of $\mathbb{R}^d$ in the first d-coordinates. A pyramid of P is the convex hull of P and any single point not in the $x_{d+1} = 0$ hyperplane. Determine the face poset of a pyramid of P in terms of $\mathcal{F}(P)$. A bipyramid of P is the convex hull of P and two points on opposite sides of the hyperplane $x_{d+1} = 0$ whose line segment goes through the relative interior of P. What is the h-vector of a bipyramid of P?*

Given any poset $(L, \leq)$ we can define the *dual* poset $L^\star$ by reversing the inequality sign. Formally, $x \leq_{L^\star} y$ if and only if $y \leq_L x$. Face posets of polytopes have the following remarkable property: $(\mathcal{F}(P))^\star$ is isomorphic to the face poset of some (usually different) polytope.

7. Polarity

One of the most useful way to construct polytopes is the polar or dual polytope. Let $(\mathbb{R}^d)^\star$ be the vector space of all linear functionals, i.e., linear maps $l : \mathbb{R}^d \to \mathbb{R}$. The space of linear functionals is a d-dimensional vector space, so it makes sense to talk about convexity in $(\mathbb{R}^d)^\star$. If fact, $\mathbf{v} \to l_{\mathbf{v}}(\mathbf{x})$, where $l_{\mathbf{v}}(\mathbf{x}) = \mathbf{v} \cdot \mathbf{x}$ is an isomorphism from R^d to $(\mathbb{R}^d)^\star$. There is a *natural* isomorphism between $\mathbb{R}^d$ and $((\mathbb{R}^d)^\star)^\star = (\mathbb{R}^d)^{\star\star}$. It is given by "evaluation". Given $\mathbf{x} \in \mathbb{R}^d$, define $\mathbf{x}^{\star\star} : (\mathbb{R}^d)^\star \to \mathbb{R}$ by $\mathbf{x}^{\star\star}(l) = l(\mathbf{x})$ for any $l \in (\mathbb{R}^d)^\star$. Looking at the expression $\mathbf{v} \cdot \mathbf{x}$ we can either imagine fixing $\mathbf{v}$ and allowing $\mathbf{x} \in \mathbb{R}^d$ to vary, giving an element of $(\mathbb{R}^d)^\star$, or we can imagine fixing $\mathbf{x}$, and letting $\mathbf{v} \in (\mathbb{R}^d)^\star$ vary, giving an element of $(\mathbb{R}^d)^{\star\star}$. In either case, the equation $\mathbf{v} \cdot \mathbf{x} = c$ is a hyperplane in the corresponding space.

Example 7.1. *Let $\mathbf{v} = (1, 2, 3, 4) \in (\mathbb{R}^4)^\star$ and $\mathbf{x} = (5, 6, 7, 8) \in \mathbb{R}^4$. In $\mathbb{R}^4$ the equation $\mathbf{v} \cdot \mathbf{x} = 9$ is the hyperplane $x_1 + 2x_2 + 3x_3 + 4x_4 = 9$. In $(\mathbb{R}^4)^\star$ the hyperplane is $5v_1 + 6v_2 + 7v_3 + 8v_4 = 9$.*

Let K be a closed and bounded convex subset of R^d such that $\mathbf{0}$ is in the interior of K. The *polar* set of K is

$$K^\star = \{l \in (\mathbb{R}^d)^\star : \forall \mathbf{x} \in K,\ l(\mathbf{x}) \leq 1\}.$$

Figure 19 shows that in $\mathbb{R}^2$ the polar of the square is the diamond.

Theorem 7.2. *Let K and $K^\star$ be as above.*

(1) *$K^\star$ is a closed and bounded convex set which contains $\mathbf{0}$ in its interior.*
(2) $K = K^{\star\star}$.
(3) *Let F be a face of K. Define $\Psi(F) = \{l \in K^\star : l(\mathbf{x}) = 1$ for all $\mathbf{x} \in F.\}$ Then $\Psi(\Psi(F)) = F$ and Ψ is an inclusion reversing bijection from the faces of K to the faces of $K^\star$.*

Proof. We leave 1 and the fact that $K \subseteq K^{\star\star}$ as an exercise. Suppose $\mathbf{x} \notin K$. Then there exists $l \in (\mathbb{R}^d)^\star$ such that $l(\mathbf{x}) > 1$ and $l(\mathbf{y}) \leq 1$ for all $y \in K$. By definition, such an l is in $K^\star$. Therefore, $x^{\star\star} \notin K^{\star\star}$.

It is clear that Ψ is inclusion reversing. (Right?) We still need to show that $\Psi(F)$ is a face of $K^\star$ and that $\Psi(\Psi(F)) = F$. Let $\mathbf{y}$ be in the relative interior of F. Let $F^\star = \{l \in K^\star : l(\mathbf{y}) = 1.\}$ Evidently, $\Psi(F) \subseteq F^\star$. Suppose $l_0 \in F^\star$, but $l_0 \notin \Psi(F)$. Then there exists $\mathbf{z} \in F, l_0(\mathbf{z}) < 1$. As $\mathbf{y}$ is in the relative interior of F, there exists $\mathbf{w} \in F$ such that $\mathbf{y} = \lambda \mathbf{z} + (1-\lambda)\mathbf{w}, 0 < \lambda < 1$. This means that $\mathbf{y}$ is on the line segment in F from $\mathbf{z}$ to $\mathbf{w}$. But $l_0(\mathbf{z}) < 1$ and $l_0(\mathbf{y}) = 1$, so $l_0(\mathbf{w}) > 1$, contrary to the choice of $l_0 \in F^\star \subseteq K^\star$.

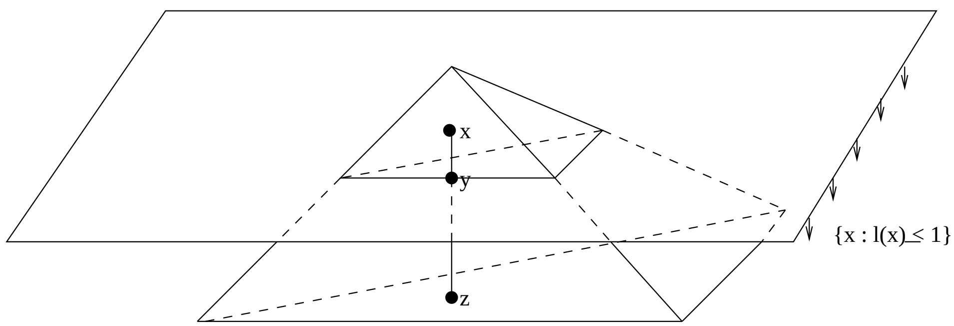

FIGURE 20

Finally, $\Psi(\Psi(F)) = \{\mathbf{x}^{\star\star} : l(\mathbf{x}) = 1$ for all $l \in \Psi(F).\}$ Hence, $F \subseteq \Psi(\Psi(F))$. For the reverse inclusion, suppose that $\mathbf{x} \notin F$ and $\mathbf{x} \in K$. Let $l_\mathbf{v} = 1$ be a supporting hyperplane of F. As $\mathbf{x} \in K - F, l_\mathbf{v}(\mathbf{x}) < 1$. Since $l_\mathbf{v} = 1$ is a supporting hyperplane of F, $l_\mathbf{v} \in \Psi(F)$. Therefore, $\mathbf{x}^{\star\star} \notin \Psi(\Psi(F))$. ☐

When P is a polytope, $P^\star$ is known as the dual polytope. Why is it a polytope? What are its vertices? In what sense are the $\mathcal{H}$ and $\mathcal{V}$-polytope definitions dual? What is the f-vector of $P^\star$? Can you see that $\mathcal{F}(P^\star)$ is isomorphic (as a poset) to $(\mathcal{F}(P))^\star$?

The dual of a simplicial polytope is called a *simple* polytope. It has the property that every vertex is incident to exactly d edges, where d is the dimension of the polytope. (Why?) Conversely, the dual of any d-polytope all of whose vertices are incident to exactly d edges must be a simplicial polytope. (Why?)

Problem 7.3. *What, if any, is the relationship between the face poset of the dual of a pyramid of P and the face poset of a pyramid of the dual of P?*

8. Zonotopes and Hyperplane arrangements

We have seen that any polytope is the affine projection of a simplex. Other properties of polytopes can be described in terms of affine projections.

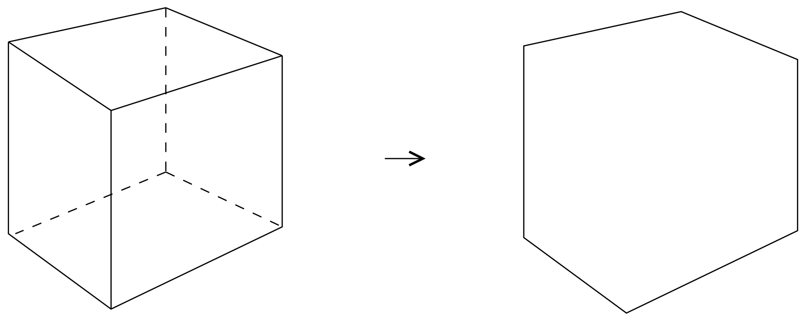

FIGURE 21. Hexagon as a projection of a cube

Definition 8.1. *A polytope P is centrally symmetric if $\mathbf{x} \in P$ if and only if $-\mathbf{x} \in P$.*

Cubes and cross polytopes are certainly centrally symmetric.

Problem 8.2. *Show (directly) that P is centrally symmetric if and only if P is an affine (in fact linear) image of a cross polytope. Show that P is centrally symmetric if and only if $P^\star$ is centrally symmetric.*

What about cubes?

Definition 8.3. *A zonotope is any affine projection of the cube.*

Any zonotope is just a translation away from being centrally symmetric, so we will assume that this is the case. Zonotopes can also be described in terms of Minkowski sums.

Definition 8.4. *Let P, Q be subsets of $\mathbb{R}^d$. Their Minkowski sum is*

$$P + Q = \{\mathbf{x} + \mathbf{y} : \mathbf{x} \in P, \mathbf{y} \in Q\}.$$

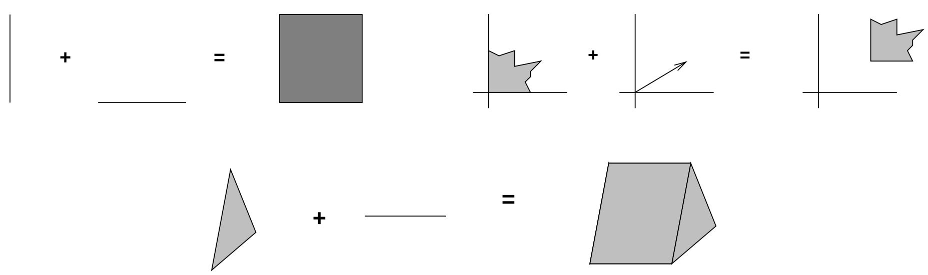

FIGURE 22. Cube, translation and prism as a Minkowski sum

In order to analyze the connection between zonotopes and hyperplane arrangements we must first discuss fans and cones.

Problem 8.5. *If P and Q are polytopes in $\mathbb{R}^d$ is $P + Q$ a polytope? Show that a P is a zonotope if and if it is a translation of the Minkowski sum of line segments of the form $[-\mathbf{x}_i, \mathbf{x}_i]$.*

Definition 8.6. *A subset C of $\mathbb{R}^d$ is a cone if for all $\mathbf{x} \in C$ and $t \geq 0, t\mathbf{x} \in C$.*

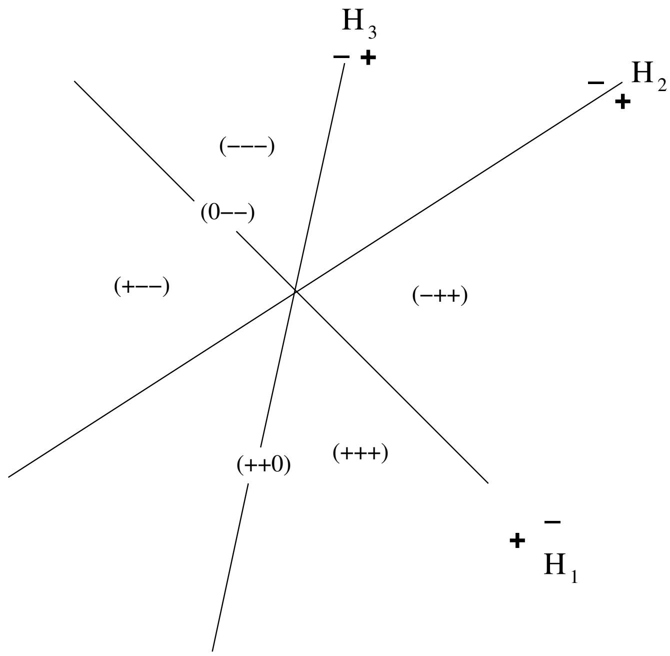

FIGURE 23. Several sign vectors

For any $A \subseteq \mathbb{R}^d$ we can form the *conical hull* of A, the smallest cone which contains A. You can check that the conical hull of A is $\{t\mathbf{x} : \mathbf{x} \in A, t \geq 0\}$.

Let $\mathcal{A} = \{H_1, \ldots, H_n\}$ be a collection of linear hyperplanes (the origin is in every hyperplane) in $\mathbb{R}^d$. For each H_i let $\mathbf{v}_i$ be one of the two unit vectors normal to H_i. Choosing $\mathbf{v}_i$ is equivalent to deciding which half of the complement of H_i will be designated positive. There is a natural set of cones associated with the $\mathbf{v}_i$. For every $\mathbf{x} \in \mathbb{R}^d$ we associate a sign vector $s(\mathbf{x}) \in \{+, 0, -\}^n$ depending on whether $\mathbf{v}_i \cdot \mathbf{x}$ is positive, zero or negative. The set of all $\{s(\mathbf{x}) \in \{+, 0, -\}^n : \mathbf{x} \in \mathbb{R}^d\}$ is called the set of *covectors* of $\mathcal{A}$ and encodes a great deal of the information about the combinatorics of $\mathcal{A}$. Figure 23 shows several sign vectors of an arrangement of three lines in R^2.

If we fix a sign vector $s \in \{+, 0, -\}^n$ then we define $C(s) = \{\mathbf{x} \in \mathbb{R}^d : \mathbf{v}_i \cdot \mathbf{x} \geq 0$ if $s_i = +, \mathbf{v}_i \cdot \mathbf{x} \leq 0$ if $s_i = -,$ and $\mathbf{v}_i \cdot \mathbf{x} = 0$ if $s_i = 0.\}$ For every sign vector s, $C(s)$ is a polyhedral cone. Of course, there may be no vectors with a particular sign vector s. In that case, $C(s) = \{\mathbf{0}\} = C(\{0, \ldots, 0\})$. The cones $C(s)$ are usually called the faces of the arrangement. The collection of all of the cones $C(s)$ is a *complete fan* in $\mathbb{R}^d$.

Definition 8.7. *A collection $Fn = \{C_1, \ldots, C_m\}$ of cones in $\mathbb{R}^d$ is a* **complete fan** *if,*

- $\cup C_i = \mathbb{R}^d$.
- *Every cone in Fn is a polyhedron.*
- *If F is a face of C_i, then F is also in Fn.*
- *If C_1 and C_2 are in Fn, then $C_1 \cap C_2$ is in Fn.*

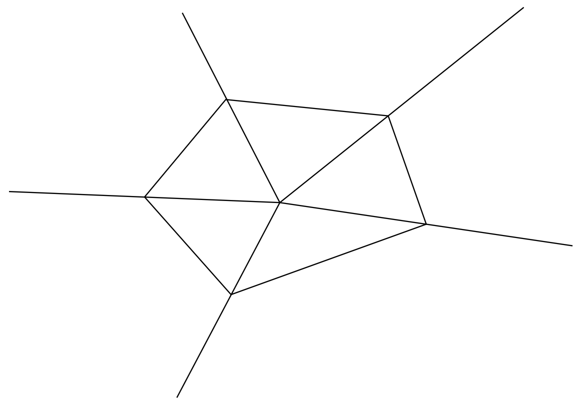

FIGURE 24. Face fan of a polygon

Normally a fan is defined as above without the first condition and a complete fan is a fan which satisfies all four conditions. As we will only consider complete fans, we will just say fan for any collection of cones which satisfy all of the above conditions. The covectors of $\mathcal{A}$ determine the face lattice of the fan of the hyperplane arrangement. Indeed, $C(s)$ is a face of $C(t)$ if and only if s can be obtained from t by changing some number of signs to zeros.

Another example of a fan is the face fan of a polytope. Suppose P is a d-polytope in $\mathbb{R}^d$ with the origin it its interior. For each face other than P, let $C(F)$ be the conical hull of F. The collection of all such cones (where we set the connical hull of $\emptyset$ equal to $\mathbf{0}$) is a fan.

The connection between hyperplane arrangements and zonotopes is best seen through the *normal fan* of a polytope. One point of view of the normal fan is through the lens of linear programming. The goal is to maximize the value of an objective function on a polytope. An objective function is always a linear functional $l_{\mathbf{v}}$. The maximum (or minimum) of such a function will always occur along faces of P. See Figure 25.

For each face of P other than $\emptyset$, let $N_F = \{\mathbf{v} \in (\mathbb{R}^d)^\star : F \subseteq \{\mathbf{x} \in P : \mathbf{v} \cdot \mathbf{x} = \max \mathbf{v} \cdot \mathbf{y}, \mathbf{y} \in P\}\}$. Each N_F is a polyhedral cone and all together they form the normal fan of P.

Proposition 8.8. *Suppose that P has $\mathbf{0}$ in its interior. Then the face fan of P is the normal fan of $P^\star$ and the normal fan of P is the face fan of $P^\star$.*

Proof. By duality it is sufficient to prove that the face fan of P is the normal fan of $P^\star$. To see that the face fan of P is contained it the normal fan of $P^\star$ we show that for $\mathbf{x} \in F, F$ a face of $P, \mathbf{x} \in N_{\Psi(F)}(P^\star)$. (Why is this enough?) But this is immediate - $\Psi(F)$ is exactly the elements l of $P^\star$ for which $l(\mathbf{x}) = 1$ on F and less than one on $P - F$.

We leave the other inclusion as an exercise.

$\square$

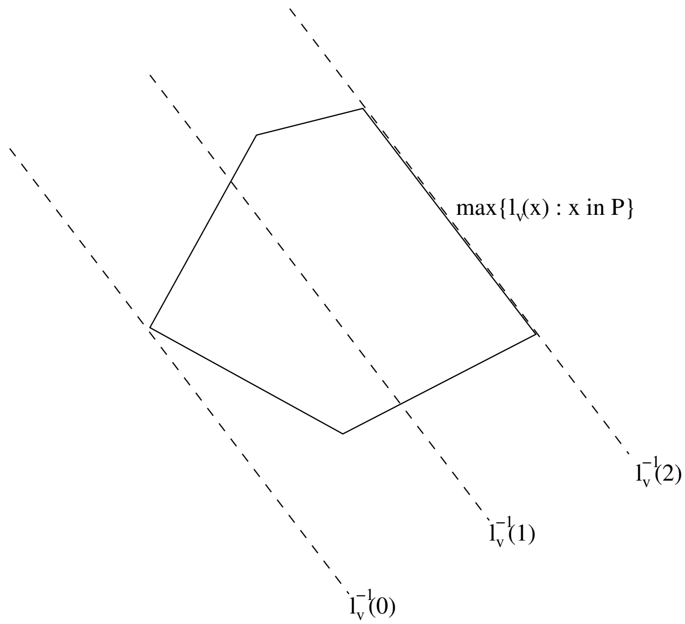

FIGURE 25. Maximizing l_v on P.

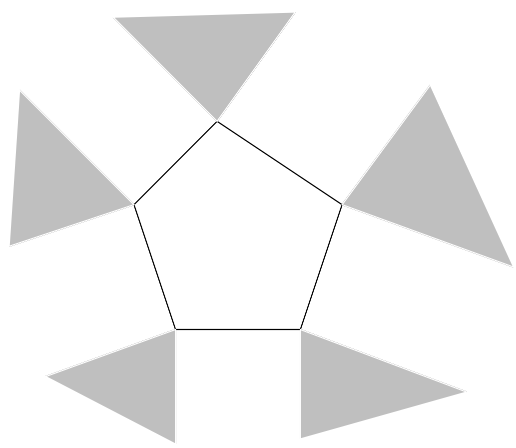

FIGURE 26. Preparing the normal fan of a pentagon

Let Z be a zonotope. We know that Z is a Minkowski sum of intervals $[-\mathbf{v}_i, \mathbf{v}_i]$. Let H_i be the corresponding perpendicular hyperplanes and let $\mathcal{A}$ be the arrangement consisting of these hyperplanes. An arrangement of hyperplanes is *essential* if $\cap H_i = \{\mathbf{0}\}$. This is equivalent to Z containing $\mathbf{0}$ in its interior.

Theorem 8.9. *Let Z and $\mathcal{A}$ be as above. Suppose that $\mathcal{A}$ is essential. Then the normal fan of Z is the fan of the hyperplane arrangement $\mathcal{A}$.*

Proof. Fix $\mathbf{y}$ and suppose that you want to maximize $\mathbf{y} \cdot \mathbf{x}$ over all $\mathbf{x} \in Z$. This is the same as maximizing over linear combinations

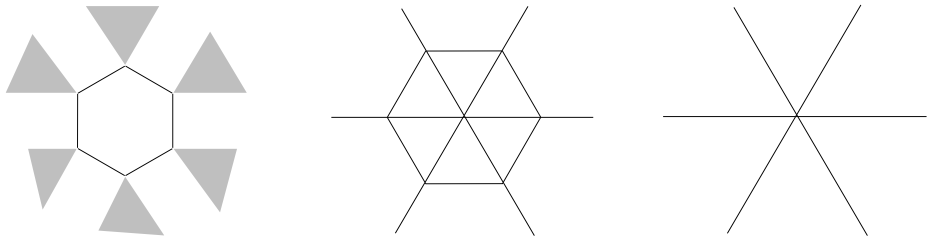

FIGURE 27. The normal fan of a hexagon, the face fan of its dual and the corresponding arrangement.

$$\mathbf{y} \cdot \{\sum_{i=1}^{n} \lambda_i \mathbf{v}_i, -1 \leq \lambda_i \leq 1\}.$$

We can do this over each summand. This achieves a maximum on the set

$$Z^{\mathbf{y}} = \sum_{i=1}^{n} \lambda_i \mathbf{v}_i \in Z : \begin{cases} \lambda_i = 1 & \text{if } \mathbf{y} \cdot \mathbf{v}_i > 0 \\ \lambda_i = -1 & \text{if } \mathbf{y} \cdot \mathbf{v}_i < 0 \\ -1 \leq \lambda_i \leq 1 & \text{if } \mathbf{y} \cdot \mathbf{v}_i = 0. \end{cases}$$

But this is the same as having the corresponding sign vector in $\mathcal{A}$.

□

The Euler characteristic of an arrangement $\mathcal{A}$ is defined in the same way it was defined for polytopes. Now $f_i(\mathcal{A})$ is the number of cones of dimension i in the fan of the arrangement.

Theorem 8.10. *If $\mathcal{A}$ is an essential arrangement of hyperplanes in $\mathbb{R}^d$, then $\chi(\mathcal{A}) = (-1)^d$.*

Proof. The fan of the arrangement is the normal fan of Z, where Z is the zonotope associated to $\mathcal{A}$. By Proposition 8.8 this is the face fan of $Z^\star$. Therefore, $f_i(\mathcal{A}) = f_{i-1}(Z^\star)$ for $i \geq 1$. Also, $f_0(\mathcal{A}) = f_d(Z^\star) = 1$. Hence, when d is even, $\chi(\mathcal{A}) = -\chi(Z^\star) + 2 = 1$. When d is odd, $\chi(\mathcal{A}) = -\chi(Z^\star) = -1$. □

9. Introduction to the Möbius function

In the next section we will derive a formula for $f_i(\mathcal{A})$, where $\mathcal{A}$ is an essential arrangement and all we know is the dimension of the intersection of any subset of hyperplanes in $\mathcal{A}$. Our main tool is the Möbius function of a poset.

Definition 9.1. *Let L be a finite poset. The* **Möbius** *function of L is the unique function $\mu : L \times L \to \mathbb{Z}$ such that*

- *If $x \not\leq y$, then $\mu(x, y) = 0$.*
- *$\mu(x, x) = 1$.*
- *If $x < y$, then $\displaystyle\sum_{x \leq z \leq y} \mu(x, z) = 0$.*

Using the above definition it is possible to determine $\mu(x, y)$ recursively for any pair $(x, y) \in L \times L$. For instance, Figure 28 shows $\mu(x, z)$ for a fixed x.

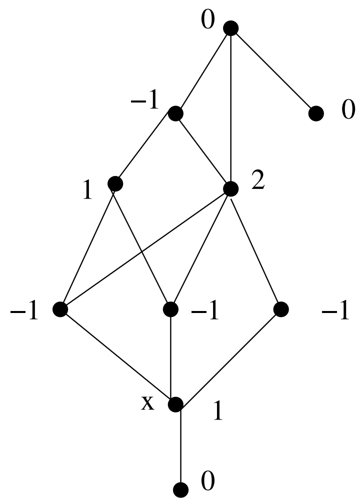

FIGURE 28. $\mu(x, z)$

Frequently it is not easy to find "nice" formulas for $\mu(x, y)$ in L. The Boolean algebra is an important exception. The poset of all subsets of $[n] = \{1, 2, \ldots, n\}$ is known as the Boolean algebra and we denote it by B_n. By definition, if $S \not\subseteq T$, then $\mu(S, T) = 0$ in B_n.

Proposition 9.2. *Let $S, T \in B_n$. If $S \subseteq T$, then $\mu(S, T) = (-1)^{|T-S|}$.*

Proof. Since $[S, T] \cong B_m$, where $m = |T - S|$, the proposition is equivalent to proving that $\mu(\emptyset, [n]) = (-1)^n$ in B_n. We do this by induction on n, with $n = 1$ being obvious. So, assume that the proposition holds for $n - 1$. By the definition of μ, the induction hypothesis, and the fact that $[\emptyset, S] \cong [\emptyset, [|S|]]$,

$$\mu(\emptyset, [n]) = -\sum_{S \subset [n]} \mu(\emptyset, S) = -\sum_{i=1}^{n-1} \binom{n}{i} (-1)^i = (-1)^n.$$

The last equality comes from the binomial formula for $(x - 1)^n$ with $x = 1$. □

Another poset for which we can easily compute μ is the poset $\mathbb{Z}^+$ with binary relation $a \leq b$ if and only if $a|b$. Now, $\mathbb{Z}^+$ is not finite. However, $\mu(a, b)$ only depends on the finite number of elements in $[a, b]$. (Remember, this interval is not the same as $[a, b]$ in $\mathbb{Z}^+$ with the usual order!) Once we realize that $[a, b] \cong [1, b/a]$ we only have to compute $\mu(1, n)$ for all n.

Suppose that n is a product of distinct primes, $n = p_1 \cdot p_2 \cdots p_m$. Then $[1, n] \cong B_m$. So, $\mu(1, n) = (-1)^m$. What if the prime factorization of n contains a square (or higher) power of a prime?

Problem 9.3. *Prove that if the prime factorization of n contains a square (or higher) power of a prime, then $\mu(1, n) = 0$.*

The function $\mu(n) = (-1)^m$, when n is a product of m distinct primes, and 0 otherwise, has been known as the Möbius function in classical number theory since the middle of the nineteenth century.

The main value of the Möbius function of a poset is the Möbius inversion formula. Before stating it we introduce a useful convolution algebra. Let $I(L) = \{\alpha : L \times L \to \mathbb{Z}, \alpha(x,y) = 0 \text{ if } x \not\leq y.\}$ Define

$$(\alpha \circ \beta)(x,y) = \sum_{x \leq z \leq y} \alpha(x,z) \cdot \beta(z,y).$$

If $x \not\leq y$, then $(\alpha \circ \beta)(x,y) = 0$. Let δ be the usual Dirac delta function and define $\zeta(x,y) = 1$ if $x \leq y, 0$ otherwise.

Lemma 9.4. *Let $\alpha, \beta, \gamma \in I(L)$.*

(1) $(\alpha \circ \beta) \circ \gamma = \alpha \circ (\beta \circ \gamma)$.
(2) $\delta \circ \alpha = \alpha \circ \delta = \alpha$.
(3) $(\zeta \circ \alpha)(x,y) = \sum_{x \leq z \leq y} \alpha(z,y)$ *and* $(\alpha \circ \zeta)(x,y) = \sum_{x \leq z \leq y} \alpha(x,z)$
(4) $\mu \circ \zeta = \delta = \zeta \circ \mu$.

Proof. For (1) we note that both sides of the equation are equal to

$$\sum_{x \leq y \leq z \leq w} \alpha(x,y)\beta(y,z)\gamma(z,w).$$

Each of (2),(3) and the left-half of (4) are immediate consequences of the definitions. We leave the right-hand side of (4) as an exercise. □

Given a function $f : L \to \mathbb{Z}$ and $\alpha \in I(L)$ define $f\alpha(y) = \sum_{x \leq y} f(x)\alpha(x,y)$. It is easy to check that $(f\alpha)\beta = f(\alpha \circ \beta)$.

Theorem 9.5 (Möbius inversion)**.** *Let f and g be functions from L to $\mathbb{Z}$. Then,*

$$g(y) = \sum_{x \leq y} f(x) \leftrightarrow f(y) = \sum_{x \leq y} g(x)\mu(x,y).$$

Proof.

$$g(y) = \sum_{x \leq y} f(x) \leftrightarrow g = f\zeta \leftrightarrow g\mu = f(\zeta \circ \mu) = f\delta = f.$$

□

If we define $\alpha f(x) = \sum_{x \leq y} \alpha(x,y) f(y)$, then the same reasoning as above proves a dual version of the above formula.

Theorem 9.6. *Let f and g be functions from L to $\mathbb{Z}$. Then,*

$$g(x) = \sum_{x \leq y} f(y) \leftrightarrow f(x) = \sum_{x \leq y} g(y)\mu(x,y).$$

Here is a typical use of Möbius inversion. Let $\phi(n)$ be the number of integers $i, 1 \leq i \leq n$ such that n and i are relatively prime. So, $\phi(6) = 2$ and $\phi(p) = p - 1$ for any prime p.

Proposition 9.7.

$$\sum_{a|n} \phi(a) = n.$$

Proof. Write down the fractions $1/n, 1/n, \ldots, (n-1)/n, n/n$. Now reduce each one to lowest terms. Each $a|n$ will appear in the denominator exactly $\phi(a)$ times. □

As noted above, in number theory $\mu(n) \equiv \mu(1,n)$ in the poset $(\mathbb{Z}^+, |)$.

Corollary 9.8.
$$\sum_{a|n} a \cdot \mu(n/a) = \phi(n).$$

Proof. Apply the last proposition and Möbius inversion to the poset $(\mathbb{Z}^+, |)$ $f(a) = \phi(a)$ and $g(n) = n$. □

10. Characteristic polynomial and hyperplane arrangements

The other ingredient for determining the number of faces in every dimension in the fan of a hyperplane arrangement is the *intersection poset* of the arrangement.

Definition 10.1. *Let $\mathcal{A} = \{H_1, \ldots, H_n\}$ be an arrangement of hyperplanes in $\mathbb{R}^d$. The* **intersection poset** *of $\mathcal{A}$, denoted by $L(\mathcal{A})$, is the poset of intersections of members of $\mathcal{A}$ ordered by* **reverse** *inclusion. The whole space (equal to the empty intersection) is always the least element of $L(\mathcal{A})$. If $\mathcal{A}$ is an affine arrangement, then $\emptyset$ is never in $L(\mathcal{A})$.*

The elements of $L(\mathcal{A})$ are called the *flats* of the arrangement. The intersection poset is always a ranked poset. The rank of $X \in L(\mathcal{A})$ is always $d - \dim X$. Figure 29 shows a simple example of $L(\mathcal{A})$.

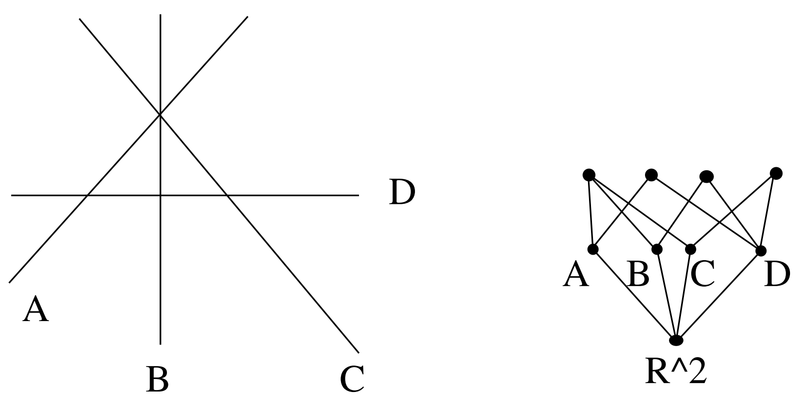

Figure 29. An example of $L(\mathcal{A})$

An arrangement is *central* if the intersection of all the hyperplanes is non-empty. Equivalently, $L(\mathcal{A})$ has a maximal element. A central arrangement is *essential* if this intersection is a point. All non-central arrangements are considered to be essential. Every arrangement of linear hyperplanes is central and an affine arrangement is central if and only if it is a translation of a linear arrangement.

Problem 10.2. *For each $i \in [n]$ let $\mathbf{v}_i$ be a non-zero vector orthogonal to H_i and let $E = \{\mathbf{v}_1, \ldots, \mathbf{v}_n\}$. Define $L'(\mathcal{A}) = \{E \cap W : W$ is a subspace of$\mathbb{R}^d.\}$ Order $L'(\mathcal{A})$ by inclusion. What is the relationship between $L(\mathcal{A})$ and $L'(\mathcal{A})$?*

For any ranked poset with a minimum element $\hat{0}$, we can collect information about the Möbius function $\mu(\hat{0}, X)$ sorted by rank as follows.

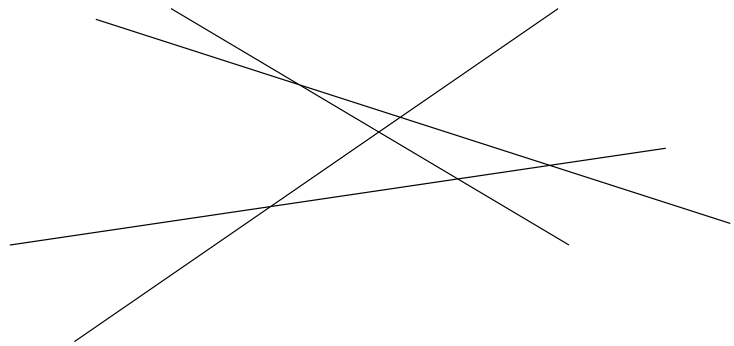

FIGURE 30. Four affine lines in general position, $\chi_{L(\mathcal{A})} = \lambda^2 - 4\lambda + 6$.

Definition 10.3. *Let L be a finite ranked poset with minimum element $\hat{0}$ with rank$(L) = d$. The* **characteristic polynomial** *of L is*

$$\chi_L(\lambda) = \sum_{X \in L} \mu(\hat{0}, X)\lambda^{d-\mathrm{rk}(X)}.$$

Example 10.4. *Let $\mathcal{A}$ be four (linear) planes in $\mathbb{R}^3$ in general position. Then $\chi_{L(\mathcal{A})} = \lambda^3 - 4\lambda^2 + 6\lambda - 3$.*

Theorem 10.5. *Let $\mathcal{A}$ be an arrangement in $\mathbb{R}^d$. Then the number of $d-i$-dimensional faces in the fan of the arrangement is*

$$\sum_{\mathrm{rk}\, X = i, Y \geq X} \mu(X, Y)(-1)^{d-\mathrm{rk}(X)}.$$

So, if $\mathcal{A}$ is essential, then the number of d-dimensional faces is $(-1)^d\chi_L(\mathcal{A})(-1)$.

Proof. Each $X \in L(\mathcal{A})$ is an affine subspace of $\mathbb{R}^d$. Each hyperplane of $\mathcal{A}$ either contains X or intersects X in a hyperplane of X. So, we can define $\mathcal{A}^X$ to be the hyperplane arrangement on X consisting of all the intersections $H_j \cap X$, where $X \not\subseteq H_j$. Let $\sigma(X)$ be the number of top-dimensional faces in $\mathcal{A}^X$. Now define $f(X) = (-1)^{\dim X}\sigma(X) = (-1)^{d-\mathrm{rk}\, X}\sigma(X)$. Every $(d-i)$-dimensional face of the fan of the arrangement occurs as a top-dimensional face in exactly one flat (its affine span). Hence, Euler's formula tells us that

$$\sum_{Y \geq X} f(Y) = (-1)^{\dim X} = (-1)^{d-\mathrm{rk}\, X}.$$

The theorem now follows from Möbius inversion.

□

It is not difficult to reconstruct $L(\mathcal{A})$ just from the knowledge of the dimension of the intersection of any subset of hyperplanes in $\mathcal{A}$. The above formula is originially due to Zaslavsky [23]. Although his proofs were not published until five years later [12], at the same time Las Vergnas gave a formula which is logically equivalent to Zaslavsky's [11]. However, neither was the first to prove that the number of faces in the top dimension only depended on the combinatorics of the intersection poset. This was due to Winder [22] - a bit of history that is still virtually unknown to many.

Problem 10.6. *What is the maximum number of d-dimensional faces that an arrangement with n hyperplanes in $\mathbb{R}^d$ can have?*

Problem 10.7. *What is the fewest number of d-dimensional faces that an essential central arrangement with n hyperplanes in $\mathbb{R}^d$ can have?*

For affine arrangements we could count only the bounded regions. One way to obtain affine arrangements is through central arrangements. Given a central arrangement, choose one hyperplane to be the "hyperplane at infinity". This leaves the remaining hyperplanes as an affine arrangement. An equivalent point of view is to intersect the arrangement with the unit sphere. This partitions the sphere into various cells corresponding to the faces of the fan of the arrangement. Now cut the sphere along any of the equators and flatten it out to obtain an affine arrangement. Figure 31 shows that choosing different hyperplanes at infinity may lead to different affine arrangements. Nonetheless......

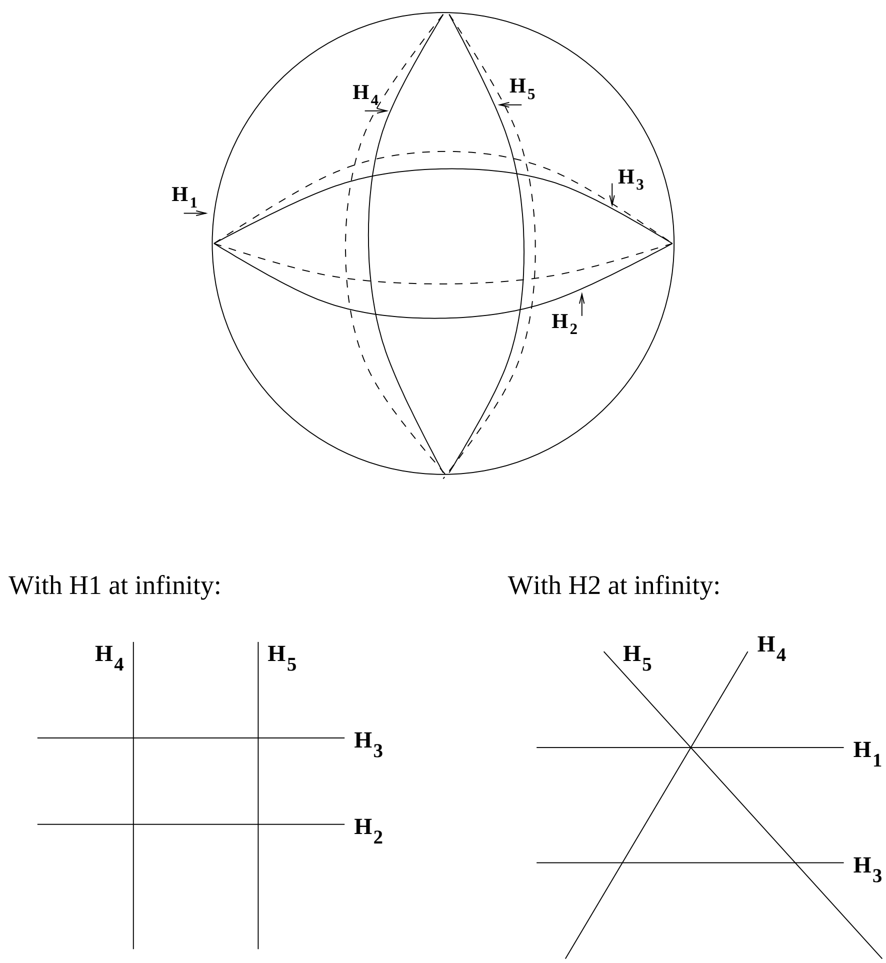

FIGURE 31. Two distinct affine arrangements associated to five hyperplanes in $\mathbb{R}^3$.

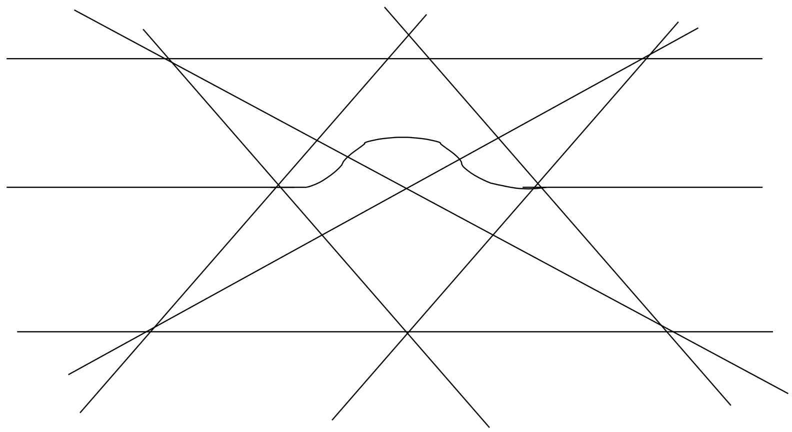

FIGURE 32. Non-Pappus pseudoline arrangement

Theorem 10.8. *Let $\mathcal{A}$ be an essential central arrangement in $\mathbb{R}^d$. Let $\mathcal{A}'$ be any affine arrangement obtained by choosing one of the hyperplanes to be the hyperplane at infinity. Then the number of d-dimensional bounded regions in $\mathcal{A}'$ is*

$$\chi'_{L(\mathcal{A})}(\lambda)|_{\lambda=1}.$$

Like the previous theorem, this was proved simultaneously by Zaslavsky and Las Vergnas.

Example 10.9. *The characteristic polynomial of the above arrangement is $\lambda^3 - 5\lambda^2 + 8\lambda - 4$. Direct computation shows that $\chi'_{L(\mathcal{A})}(\lambda)|_{\lambda=1} = 1$.*

The results of this section can be generalized in several directions. One direction is to consider complex hyperplane arrangements. Now the complement of such an arrangement is connected, hence it does not make sense to start counting regions. However, the coefficients (or more precisely their absolute value) of the characteristic polynomial of the intersection poset have a natural topological interpretation. They are the dimension of a cohomology group of the complement.

Another direction is to allow "wavy" hyperplanes in $\mathbb{R}^d$. By putting reasonable restrictions on these pseudo-hyperplanes and how they intersect we can obtain combinatorically new arrangements, all of whom still satisfy the last two formulas. An example of this is the non-Pappus pseudoline arrangement in Figure 32.

11. FROM HYPERPLANES TO GRAPHS

Let G be a connected simple graph with vertices $\{v_1, \ldots, v_d\}$. Recall in the introduction that an acyclic orientation of G is an assignment of a direction to each edge so that there are no directed circuits. We label an edge between v_i and v_j by e_{ij}. For each e_{ij} let H_{ij} be the hyperplane $x_i - x_j = 0$ in $\mathbb{R}^d$. We can think of H_{ij} as saying that the i^{th} and j^{th} coordinates are the same. The *graph arrangement* associated to G is $\mathcal{A}_G = \{H_{ij} : e_{ij}$ an edge in $G.\}$

Proposition 11.1. *The number of d-dimensional regions in the hyperplane fan for $\mathcal{A}_G$ equals the number of acyclic orientations of G.*

Proof. Let $\mathbf{x}$ be in the interior of a d-dimensional face F in the fan of $\mathcal{A}$. The interiors of the highest dimensional faces are in the complement of $\cup H_i$. Therefore, $\mathbf{x}_i \neq \mathbf{x}_j$ for every edge e_{ij} in G. If $\mathbf{x}_i < \mathbf{x}_j$, then orient the edge e_{ij} from v_i to v_j. Otherwise orient e_{ij} the other direction. Call this orientation O_F. This must be an acyclic orientation. As F is a cone, the orientation is independent of the choice of $\mathbf{x}$ in F. Conversely, given an acyclic orientation O of G it is not difficult to find $\mathbf{x} \in \mathbb{R}^d$ such that $x_i < x_j$ for each edge $e_{ij} \in G$ which is oriented from v_i to v_j. Such an $\mathbf{x}$ must be in a d-dimensional face of the fan. Call this face F_O. The two maps, $F \to O_F$ and $O \to F_O$ are inverses of each other. $\square$

Corollary 11.2. *The number of acyclic orientations of G is $(-1)^d \chi_{L(\mathcal{A}_G)}(-1)$.*

The above formula is originally due to Stanley [16]. As stated, it is somewhat unsatisfactory as it requires us to go through $\mathcal{A}_G$ instead of directly appealing to G. We do this in two (apparently) different ways.

Definition 11.3. *Let $\lambda \in \mathbb{Z}+$ and let G be a graph (not necessarily simple). $p_G(\lambda)$ is the number of ways to properly color G with λ colors.*

$p_G(\lambda)$ is called the *chromatic polynomial* of G. It was introduced by Birkhoff in 1912 in an attempt to solve the four-color problem.

Let e be an edge of G. The *deletion* of e from G is the graph obtained by deleting the edge. It is denoted $G - e$. The *contraction* of e in G is the graph obtained by removing e and identifying its two vertices to one. The contraction of e is denoted G/e If e is a loop, then $G - e = G/e$.

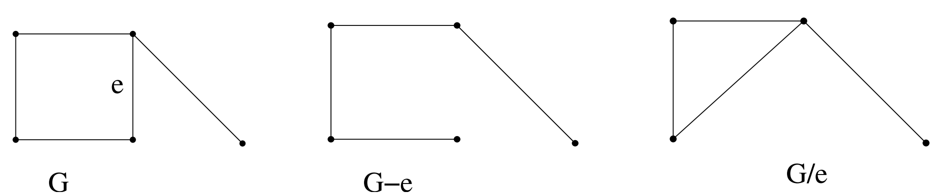

FIGURE 33. Deletion and contraction at e.

Proposition 11.4. *Let e be an edge of G. If e is a loop, then $p_G(\lambda) = 0$. Otherwise, $p_G(\lambda) = p_{G-e}(\lambda) - p_{G/e}(\lambda)$.*

Proof. Identify colorings of G/e with colorings of G whose colors on the identified vertices are the same. Then proper colorings of G correspond to proper colorings of $G - e$ which are different on the vertices of e, i.e, those which do not correspond to a proper coloring of G/e. $\square$

Corollary 11.5. *$p_G(\lambda)$ is a monic polynomial of degree equal to the number of vertices of G.*

Proof. Apply deletion/contraction and induct on the number of edges of G. $\square$

Theorem 11.6. [16] *The number of acyclic orientations of G is $(-1)^d p_G(-1)$.*

Proof. Induction on d and deletion/contraction. For e a non-loop edge, this requires us to prove that the number of acyclic orientation of G equals the number of acyclic orientation of $G - e$ plus the number of acyclic orientations of G/e. We leave this as an exercise. $\square$

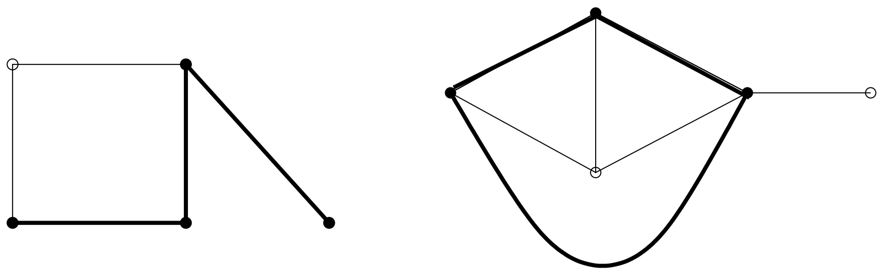

FIGURE 34. Flats of a graph

A different approach to the chromatic polynomial is via Möbius inversion. A *flat* or *closed* subgraph of G is any subgraph induced by a partition of the vertices. The poset of flats of G ordered by inclusion is a lattice of rank $d-k$, where k is the number of components of G, and is called the lattice of flats of G and we will denote it $L(G)$. In fact,.....

Problem 11.7. *Show that* $L(G) \cong L(\mathcal{A}_G)$.

Any coloring of the vertices of G induces a natural partition of G whose blocks are the vertices of a single color. Thus, each coloring is associated to a unique flat $X \in L(G)$. Proper colorings are exactly those which correspond to the minimum flat of $L(G)$. For each flat X define $f(X)_\lambda$ to be the number of colorings with λ colors associated to X. Therefore, $p_G(\lambda) = f(\hat{0})_\lambda$.

Theorem 11.8. *Let G be a simple graph with k components. Then,*

$$p_G(\lambda) = \lambda^k \chi_{L(G)}(\lambda).$$

Proof. What is $\sum_{X \le Y} f(Y)_\lambda$? Evidently, this is all possible colorings of X. Let $k(X)$ be the number of components of X. Hence, $\sum_{X \le Y} f(Y)_\lambda = \lambda^{k(X)}$. Now apply Möbius inversion to obtain

$$p_G(\lambda) = f(\hat{0})_\lambda = \sum_Y \mu_{L(G)}(\hat{0}, Y)\lambda^{k(Y)} = \sum_Y \mu_{L(G)}(\hat{0}, Y)\lambda^k \cdot \lambda^{d-\mathrm{rk}(Y)}.$$

□

Problem 11.9. $p_G(-1)$ *can be interpreted as the number of acyclic orientations of G. Can you find a combinatorial interpretation for the chromatic polynomial evaluated at other negative integers?*

Problem 11.10. *Let Π_n be the partition lattice. Members of Π_n are partitions of $[n]$ ordered by refinement. Compute* $\chi_{\Pi_n}(\lambda)$.

As observed earlier, the coefficients of $\chi_{L(G}(\lambda)$ alternate in signs. It seems logical that there should be a combinatorial interpretation for the coefficients. It is also true that the coefficients of $\chi_{L(\mathcal{A})}$ alternate in sign for any hyperplane arrangement. Is there a combinatorial interpretation for these coefficients?

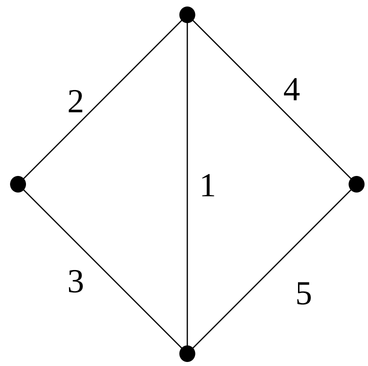

FIGURE 35

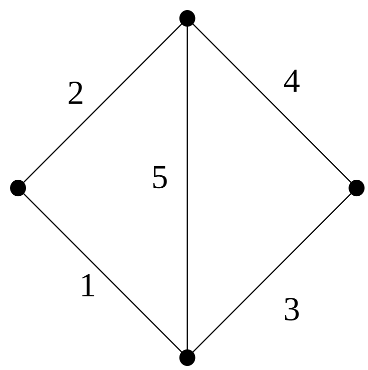

FIGURE 36. A different order

12. BROKEN CIRCUITS

The broken circuit concept was introduced by H. Whitney [20] as a method for analyzing the coefficients of $p_G(\lambda)$. In order define broken circuits we must first order the edges of G. Given such an order, a *broken circuit* of G is a circuit with its least edge removed. The broken circuits of the graph of Figure 35 are $\{2,3\}, \{4,5\}$ and $\{3,4,5\}$. Obviously, if we choose a different order for the edges we will have different broken circuits. For instance, if we reorder the previous graph as in Figure 36, the broken circuits are $\{2,5\}, \{4,5\}$ and $\{2,3,4\}$.

Since any subset of a set which does not contain a broken circuit will not contain a broken circuit, the collection of all subsets of edges of G which do not contain a broken circuit is an abstract simplicial complex. It is is called the *broken circuit* complex or the *NBC-complex.* While the broken circuit concept is due to Whitney, the idea of forming a simplicial complex out of the broken cricuit-free set is originally due to Wilf [21]. "Naturally" our first reaction is to check the f-vector of the complex. Even though the complexes for the two different orderings above are different, their f-vectors are $(1,5,8,4)$. Checking the chromatic polynomial of G we find that $p_G(\lambda) = \lambda(\lambda^3 - 5\lambda^2 + 8\lambda - 4)$. Whitney's remarkable result is that this is not just a coincidence.

Theorem 12.1. [20] *Let G be a graph with d vertices and k components. Then*

$$p_G(\lambda) = \lambda^k(f_{-1}\lambda^{d-k} - f_0\lambda^{d-k-1} + f_1\lambda^{d-k-2} - \cdots + (-1)^{d-k}f_{d-k-1}),$$

where $(f_{-1}, f_0, \ldots, f_{d-k-1})$ is the f-vector of any broken circuit complex of G.

Notice that if G has a loop, then $\emptyset$ is a broken circuit. Hence, every set (including the empty set) contains a broken circuit and all the f_i are zero.

At first sight Whitney's theorem looks very unlikely. We will give Whitney's original proof. Before this, we recall the principle of inclusion-exclusion. Let $A_1, \ldots, A_n$ be subsets (not necessarily distinct) of S. For each $J \subseteq [n]$, let $A_J = \cap_{j\in J}A_j$. Finally, let $S_i = \cup_{|J|=i}A_J$. As usual, $A_\emptyset = S$, so $S_0 = S$. Then,

$$|S - \cup_j A_j| = |S_0| - |S_1| + \cdots + (-1)^n|S_n|.$$

Problem 12.2. *Prove this using Möbius inversion on a Boolean algebra.*

Proof: [Whitney's theorem] Let $E = \{e_1, \ldots, e_n\}$ be any ordering of the edges of E. Let S be the set of all λ-colorings of G. For each $j \in [n]$, let A_j be the set of all λ-colorings of the vertices of G such that the two vertices of edge e_j have the same color. Thus, $p_G(\lambda) = |S - \cup_j A_j|$. By inclusion-exclusion this is equal to $|S_0| - |S_1| + \cdots + (-1)^n|S_n|$, where S_i is $\cup_{|J|=i}A_J$ and $A_J = \cap_{j\in J}A_j$. Now, the cardinality of A_J is $\lambda^{k(J)}$, where $k(J)$ is the number of components in the subgraph of G whose edges are $E_J = \cup_{j\in J}e_j$. How big is $|S_i|$? Who knows. So we look for some cancellation.

Let $C_1, \ldots, C_k$ be an arbitrary, but fixed, ordering of the circuits of G. For each m, let $\hat{C}_m$ be the broken circuit associated to C_m. In each E_J let l_J be the least l such that E_J contains either C_l or $\hat{C}_l$. If E_J does not contain any broken circuits then we set $l_J = 0$. All of the subsets of edges E_J for which $l_J > 0$ come in pairs. One, say E_M which contains the least element of C_{l_M} and $E_{M'}$, which is E_M with the least element of l_M removed. While we do not know $|A_M|$, we do know that it is the same as $|A_{M'}|$ since the number of components will not change when we complete a circuit. As the cardinality of M and M' differ by one, the contributions of A_M and $A_{M'}$ will cancel in the alternating sum of the S_i.

Putting these cancellations into our previous formula allows us to write

$$p_G(\lambda) = |S'_0| - |S'_1| + \cdots + (-1)^n|S'_n|, \tag{3}$$

where S'_i is the union of all A_J such that $|J| = i$ and E_J does not contain a broken circuit. What is $|S'_i|$? Consider A_J with $|J| = i$. Certainly, A_J does not contain any circuits since it does not contain any broken circuits. An elementary fact from graph theory tells us that the number of components of A_J is $d-i$. (If you have not seen this try and prove it.) Therefore, $|A_J| = \lambda^{d-i}$ and $|S'| = f_{i-1}\lambda^{d-i}$. Since any subset with more that $d-k$ edges contains a broken circuit,$S'_i = \emptyset$ when $i > d-k$. Hence, we are done. □

All of these ideas can be generalized to the situation in Problem 10.2 and even more generally to what are known as geometric lattices. Sticking to the situation in Problem 10.2, let $E = \{\mathbf{v}_1, \ldots, \mathbf{v}_n\}$ be vectors in a vector space V. Define $L(E) = \{E \cap W : W \text{ a subspace of } V\}$. Then $L(E)$ is a lattice of rank d, where d is the dimension of $< E >$ and $\chi_L(E)(\lambda)$ is a monic degree d polynomial whose coefficients alternate in sign. Suitably interpreted, Whitney's theorem also holds in this situation. A circuit of $L(E)$ is a minimal dependent set in E. After ordering

E, a broken circuit is a circuit with its least element removed. The broken circuit complex of $L(E)$ consists of all subsets of E which do not contain a broken circuit. Once again, the f-vector of the broken circuit complex determines the coefficients of the characteristic polynomial $\chi_{L(E)}(\lambda)$.

As mentioned earlier, Wilf introduced the broken circuit complex in an effort to understand the coefficients of $p_G(\lambda)$. The title of this paper is,"What polynomials are chromatic?" The fact that the coefficients of the chromatic polynomial must be the f-vector of a simplicial complex already restricts the possibilities. These restrictions are known as the Kruskal-Katona theorem [9], [10]. The name is for the usual reason, Schützenberger [14] was the first to prove it. Surprisingly, (or maybe not) broken circuit complexes are always shellable. This implies that if we transform the f-vector to the h-vector using our previous prescription then the h-vector must be non-negative. Now, this had been known for a long time via deletion/contraction and you can try and prove it yourself, but the fact that $h_{i+1} \leq h_i^{<i>}$ was new and followed from recent work of Stanley [17]. Despite a vast literature concerning chromatic polynomials of graphs these are almost all of the completely known *general* facts concerning the coefficients of $p_G(\lambda)$. There are many unsolved problems concerning these invariants.

13. Problems and connections

The characteristic polynomial $\chi_{L(E)}(\lambda)$ and hence the f- and h-vectors of broken circuit complexes appear in a number of problems. Here are some examples.

(1) (Tutte) Let G be a connected graph with d vertices. Orient the edges in any fashion. A *nowhere-zero* n-flow is a function from the edges of G to $\mathbb{Z}/n\mathbb{Z} - \{0\}$ such that the sum at every vertex is zero. The number of nowhere-zero n flows does not depend on the chosen orientation. It is $\chi_{L(G)^\star}(n)$. The lattice $L(G)^\star$ is known as the dual lattice in the sense of geometric lattices. This is not the same as our previously defined dual partial order. While we will not define it here, it is of the form $L(E)$ for an easily defined set of vectors.

(2) (Crapo-Rota) Let E be a finite set of vectors in a finite vector space V over a (finite) field F, with $< E >= V$. Then the number of linear hyperplanes H such that $H \cap E = \emptyset$ is $\chi_{L(E)}(|F|)$. In fact, there is a subspace W of V of dimension k which avoids E if and only if $\chi_{L(E)}(|F|^{(d-\dim(W))}) \neq 0$, where $d = \dim V$.

(3) In addition to the complements of hyperplane arrangements we have already seen, the f- and h-vector of broken circuit complexes for $L(E)$ occur as topological invariants of quotient spaces defined by various torus actions (real or $\mathbb{Z}_p$) on spheres and other spaces.

There are remarkably few general facts known about the f- and h-vectors of broken circuit complexes. Here are some sample questions all of which are open.

(1) Is the f-vector unimodal?
(2) Is the f-vector log concave, i.e., is $f_{i-1}f_{i+1} \leq f_i^2$?
(3) Is the h-vector unimodal?
(4) Is the h-vector log concave?
(5) Let r be the greatest integer such that $h_r > 0$. Is $h_0 \leq h_1 \leq \cdots \leq h_{\lfloor r/2 \rfloor}$? Is $h_i \leq h_{r-i}$ for $i < r/2$?

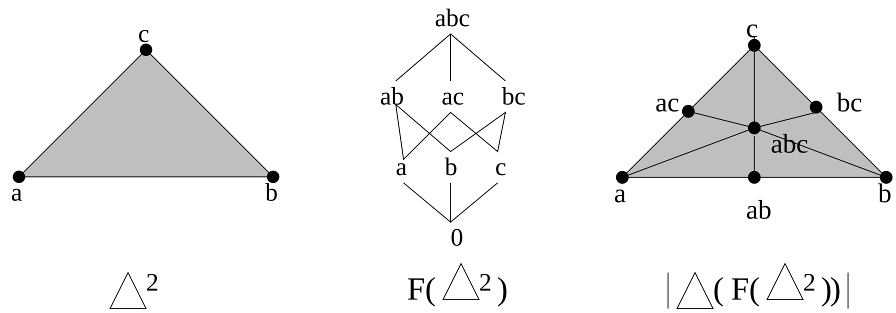

FIGURE 37. $|\mathcal{F}(\Delta^2) - \emptyset|$

Positive answers to some of these questions imply positive answers to others. (2) implies (1), and (4) implies the first three all hold. All of these problems have been open a long time and many, many special cases are known. For instance, if $\dim < E >= |E| - 2$, then the h-vector of the broken circuit complex for $L(E)$ is log concave.

Problem 13.1. *What if* $\dim < E >= |E| - 3$? *What if we restrict our attention to E a set of vectors in a vector space in a "small" field? What if we restrict our attention to broken circuit complexes of graphs with only two more edges than vertices? Or planar graphs with this property?*

We saw earlier that given a connected graph G, the h-vector of the abstract simplicial complex of edges which do not disconnect the graph directly determines the reliability of the graph. The first four questions concerning the broken circuit complex are also open for these complexes. The answer to fifth question is positive (Chari [4]) and in fact $(g_0, g_1, \dots, g_{\lfloor r/2 \rfloor})$, where $g_i = h_i - h_{i-1}$, is known to satisfy $g_{i+1} \le g_i^{<i>}$ (Hausel and Sturmfels [8], Swartz [19]). A positive answer to these questions might allow better estimates of network reliability.

Perhaps the most important question concerning f- and h-vectors was mentioned earlier. Does the conclusion of the g-theorem for simplicial polytopes hold for any simplicial complex whose geometric realization is homeomorphic to a sphere?

Recall that a simplicial complex is pure if all of its facets are the same dimension. Here is a deceptively simple question. Characterize the f-vectors of pure complexes. Not much is known here and it should be possible to improve on the following type of observation.

Proposition 13.2. *If Δ is a pure $(d-1)$-dimensional simplicial complex, then* $f_0 \le f_1 \le \cdots \le f_{\lfloor d/2 \rfloor}$.

Proof. Count incidences of the form i-simplex contained in an $(i+1)$-simplex in two different ways. □

There are strong connections between posets and topology. Face posets of polytopes and abstract simplicial complexes show that we can sometimes attach natural posets to a geometric object when it has a combinatorial structure. It is also possible to reverse this process. Let L be a finite poset. The *order complex* of L is $\Delta(L)$. It is the abstract simplicial complex whose vertices are the elements of L and whose simplices are chains in L. Figure 37 shows the order complex of $F(\Delta^2)$.

Here are three facts which demonstrate the close relationship between order complexes and geometric realizations.

Theorem 13.3. *Let Σ be any abstract simplicial complex. Then $|\Sigma|$ and $|\Delta(\mathcal{F}(\Sigma)-\emptyset)|$ are homeomorphic.*

Theorem 13.4. *Let P be a polytope. Then $|\Delta(\mathcal{F}(P)-\{\emptyset,P\})|$ is homeomorphic to ∂P.*

Theorem 13.5. *Let L be any finite poset with minimal element $\hat{0}$ and maximal element $\hat{1}$. Then $\mu(\hat{0},\hat{1})=\chi(\Delta(L-\{\hat{0},\hat{1}\}))-1$.*

A recent connection between geometry and combinatorics is the Charney-Davis conjecture. Beginning with a famous conjecture of Hopf concerning the Euler characteristic of compact nonpositively curved manifolds, Charney and Davis were led to a similar conjecture concerning cubical piecewise-linear non-positiviely curved spaces. Remarkably this conjecture turned out to be equivalent to the following.

A simplicial complex is *flag* if for every clique of edges in the complex the corresponding simplex is in the complex. Equivalently, the minimal non-faces of Δ are all of cardinality two.

Conjecture 13.6. [5] *If Δ is a flag $(2n-1)$-dimensional simplicial complex whose geometric realization is homeomorphic to a sphere, then*

$$(-1)^n[h_{2n}-h_{2n-1}+h_{2n-2}-\cdots-h_1+h_0]\geq 0.$$

This is inequality is known to hold in several special cases including the order complex of the face poset of the boundary of a polytope (Babson/Stanley).

Acknowledgement: These notes would not have been possible without Jay Schweig's tireless assistance.

References

[1] L. Billera and C. Lee. A proof of the sufficiency of McMullen's conditions for f-vectors of simplicial polytopes. *J. Comb. Theory, Ser. A*, 31:227–255, 1981.

[2] A. Björner. The unimodality conjecture for convex polytopes. *Bull. Amer. Math. Soc.*, 4:187–188, 1981.

[3] H. Bruggesser and P. Mani. Shellable decompositions of cells and spheres. *Math. Scand.*, 29:197–205, 1971.

[4] M.K. Chari. Two decompositions in topological combinatroics with applications to matroid complexes. *Trans. Amer. Math. Soc.*, 349:3925–3943, 1997.

[5] R. Charney and M. Davis. The Euler characteristic of a non-postively curved, piecewise Euclidean manifold. *Pacific J. Math.*, 171:117–137, 1995.

[6] M. Dehn. Die Eulersche formel im Zusammenhang mit dem Inhalt in der nicht-Euklidischen Geometrie. *Math. Ann.*, 61:561–568, 1905.

[7] D. Gale. Neighborly and cyclic polytopes. In V. Klee, editor, *Proc. Sympos. Pure Math, Vol. VII*, pages 225–232. A.M.S., Providence, RI, 1963.

[8] T. Hausel and B. Sturmfels. Toric hyperkähler varieties. *Documenta Math.*, 7:495–534, 2002.

[9] G. Katona. A theorem of finite sets. In *Theory of graphs (Proc. Colloq., Tihany, 1966)*, pages 187–287. Academic Press, 1968.

[10] J. Kruskal. The numer of simplices in a complex. In R. Bellman, editor, *Mathematical Optimization Techniques*, pages 251–278. Univ. of Calif. Press, 1963.

[11] M. Las Vergna. Matroïdes oreintables. *C.R. Acad. Sci. Paris, Ser. A*, 280:61–64, 1975.

[12] M. Las Vergnas. Convexity in oriented matroids. *J. Comb. Theory, Ser. B*, 25:231–243, 1980.

[13] C. Lee. *Counting the faces of simplicial convex polytopes.* PhD thesis, Cornell University, 1981.

[14] M.-P. Schützenberger. A characteristic property of certain polynomials of E. F. Moore and C. E. Shannon. In *RLE quarterly progress report no. 55*, pages 117–131. Research Lab. of Electronics, MIT, 1959.

[15] D. Sommerville. The relations connecting the angle-sums and volume of a polytope in space of n dimensions. *Proc. Royal Soc. London, Ser. A*, 115:103–119, 1927.

[16] R. Stanley. Acyclic orientations of graphs. *Discrete Math.*, 5:171–178, 1973.

[17] R.P. Stanley. Cohen-Macaulay complexes. In M. Aigner, editor, *Higher combinatorics*, pages 51–62, Dordrecht and Boston, 1977. Reidel.

[18] R.P. Stanley. The number of faces of a simplicial convex polytope. *Advances in Math.*, 35:236–238, 1980.

[19] E. Swartz. g–elements of matroid complexes. *J. Comb. Theory. Ser. B*, 88(2):369–375, 2003.

[20] H. Whitney. A logical expansion in mathematics. *Bull. of the Amer. Math. Soc.*, 38:572–579, 1932.

[21] H. Wilf. Which polynomials are chromatic? In *Theorie Combinatorie*, pages 247–256, Rome, 1973.

[22] R.O. Winder. Partitions of N-space by hyperplanes. *SIAM J. Applied Math.*, 14:811–818, 1966.

[23] T. Zaslavsky. Facing up to arrangements: face-count formulas for partitions of space by hyperplanes. *Mem. Amer. Math. Soc.*, 1(issue 1, no. 154), 1975.