Combinatorial polynomials

University of Fribourg
Spring term 2020

Disclaimer: this is a draft-in-the-making,
it will be updated during the semester.

Please report to me any correction or typos you might find
– thank you for your cooperation!

Emanuele Delucchi

September 14, 2020

Contents

Chapter 0

Some recurring notations

$\mathscr{P}(X)$ or 2^X the set of all subsets of a set X.

$[n]$ The set $\{1, 2, \ldots, n\}$. Set $[0] := \emptyset$.

B_n The boolean poset on n elements, i.e., the set $\mathscr{P}([n])$ ordered by inclusion.

$\binom{X}{k}$, where X is a finite set and $k \in \mathbb{N}$, is the set of all k-element subsets of X. If X is finite and has, say, n elements, then $|\binom{X}{k}| = \binom{n}{k}$

B^A, where A, B are sets, denotes the set of all functions $A \to B$. If A, B are finite, the cardinality of this set is $|B|^{|A|}$

Chapter 1

Graphs, colorings and flows

Disclaimer. *We will only briefly review the basics of graph theory that are strictly necessary for our purposes. For more background or further discussion of some topics the reader can refer to one of the many excellent dedicated textbooks. Our formal setup follows mostly Tutte's book [7], while we occasionally update some terminology in order to facilitate access to more modern literature such as [3, 4].*

1.1 Graphs

Definition 1.1.1. A graph $G = (V, E, h, t)$ is a quadruple consisting of a set of vertices V, a set of Edges E and two functions $h, t : E \to V$ that assign to every edge its "ends". Given any set $A \subseteq E$ of edges we let $V(A) := h(A) \cup t(A)$ be the set of all ends of edges in A.

We will often omit braces when designing one-element sets, if no need for specification arises. For instance, given $e \in E$ we will write $V(e)$ for $V(\{e\})$.

A *loop* in G is any $e \in E$ with $|V(e)| = 1$. Two edges $e, e' \in E$ are called *parallel* if $V(e) = V(e')$. The graph G is called *simple* if it has no loops nor parallel edges. A *trail* in G is any sequence $v_0, e_1, v_1, \ldots, e_k, v_k$ of vertices and edges such that $\{v_{i-1}, v_i\} = V(e_i)$ for all $i = 1, \ldots, k$. It is called *closed* (or a "cycle") if $k > 0$ and $v_0 = v_k$. A *path* is a trail where all edges and all vertices are pairwise distinct (in this case we will talk about a "path from v_0 to v_k". A *circuit* in G is a minimal closed trail, i.e., a closed trail which, after removal of any edge, is a path (in particular, every loop is a circuit).

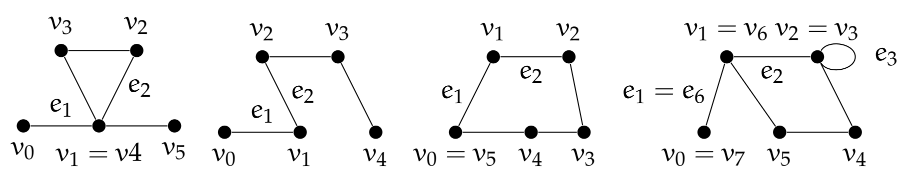

Figure 1: A trail, a path, a circuit and a closed trail that is not a circuit.

Let $T \subseteq V$ be a set of vertices of G the *vertex-induced subgraph* defined by T is the graph $G(T) := (T, E', h, t)$ where $E' = \{e \in E \mid \{h(e), t(e)\} \subseteq T\}$ is the set of edges with both endpoints in T.

Definition 1.1.2. Let G be a graph. We call G *connected* if for any two vertices $v, w \in V$ there is a path from v to w in G. A *connected component* of G is any maximal connected vertex-induced subgraph, i.e., any vertex-induced subgraph $G(T)$ that is connected and such that, for every $v \in V \setminus T$, $G(T \cup \{v\})$ is not connected . We define

$$c(G) := \text{the number of connected components of } G.$$

(a) (b)

Figure 2: Two connected graphs

Two fundamental operations on graphs are deletion and contraction of edges.

Definition 1.1.3 (Edge deletion)**.** Let $G = (V, E, h, t)$ be a graph and let $A \subseteq E$. The *deletion* of A from G is the graph $G \setminus A := (V, E \setminus A, h_{|E \setminus A}, t_{|E \setminus A})$ on the same vertex set as G but without the edges in A, and with the functions h, t restricted accordingly. If $A = \{e\}$ consists of a single element, we sometimes write $G \setminus e$ for $G \setminus \{e\}$. The "restriction" of G to A is $G[A] := G \setminus (E \setminus A)$.

Remark-Notation 1.1.4. An edge $e \in E$ is called an *isthmus* if $c(G) < c(G \setminus e)$.

Remark 1.1.5 (On the word "subgraph"). Every graph of the form $G[A]$ we will call a "subgraph" of G. Notice the difference with the notion of "vertex induced subgraph" discussed earlier on. The latter will not appear in the following, so we feel safe in our terminological choice.

Definition 1.1.6 (Edge contraction)**.** Let $G = (V, E, h, t)$ be a graph and let $A \subseteq E$. The *contraction* of A in G is the graph $G/A := (V', E', h', t')$ with edge set $E' := E \setminus A$ and vertex set $V' := V/\sim$ given as the set of equivalence classes of the equivalence relation on V generated by $v \sim w$ if $\{v, w\} = V(e)$ for some $e \in A$. The functions h' and t' are given by $h'(e) = [h(e)]_\sim$, resp. $t'(e) = [t(e)]_\sim$.

Intuitively, deleting A from G means removing every edge in A, while contracting A means "shrinking" every edge in A and identifying its endpoint-vertices.

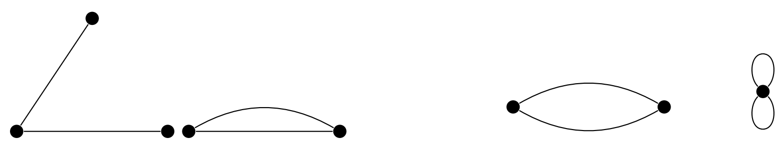

(a) (b)

Figure 3: Deletion (l.-h.s.) and contraction (r.-h.s.) of an edge in the graphs depicted in Figure 2.(a), resp. Figure 2.(b).

1.2 Colorings

Definition 1.2.1. Let $G = (V, E, h, t)$ be a graph. A *coloring* of G is any function $\gamma : V \to \mathbb{N}_{>0}$ such that $\gamma(h(e)) \neq \gamma(t(e))$ for every $e \in E$. Given $k \in \mathbb{N}$, we call γ a k-*coloring* of G if $\gamma(V) \leqslant k$. The set of all k-colorings of a given graph G is denoted by $\Gamma_G^{(k)}$.

Definition 1.2.2. Let $G = (V, E, h, t)$ be a finite graph. The *chromatic function* of G is the function

$$\chi_G^{\mathbb{N}} : \mathbb{N} \to \mathbb{N}, \quad k \mapsto |\Gamma_G^{(k)}|.$$

Example 1.2.3.

(1) If G has a loop, then obviously $\chi_G^{\mathbb{N}}(k) = 0$ for all k.

(2) If e, e' are two parallel edges of G, then $\chi_G^{\mathbb{N}}(k) = \chi_{G\setminus e}^{\mathbb{N}}(k) = \chi_{G\setminus e'}^{\mathbb{N}}(k)$ for all k.

(3) If G has no edges, i.e., $E = \emptyset$, then $\chi_G^{\mathbb{N}}(k) = k^{|V|}$ for all k.

(4) The *complete graph on* n *vertices* is the simple graph K_n on an n−element vertex set with one edge connecting any two distinct vertices. Then, $\chi_{K_n}^{\mathbb{N}}(k) = k!/(k-n)!$ for all k.

(a) $\chi_G^{\mathbb{N}}(k) = k(k-1)(k-2)^2$ (b) $\chi_G^{\mathbb{N}}(k) = k(k-1)(k-2)$

Figure 4: Two graphs with their chromatic functions.

.

Proposition 1.2.4. *Let* G *be a finite graph, and let* $e \in E$ *be any edge of* G *that is not a loop. Then,*

$$\chi_G^{\mathbb{N}}(k) = \chi_{G\setminus e}^{\mathbb{N}}(k) - \chi_{G/e}^{\mathbb{N}}(k) \textit{ for all } k \in \mathbb{N}$$

Together with Examples 1.2.4.(1) and (3), this determines the chromatic function uniquely.

Proof. We clearly have $\Gamma^{(k)}_G \subseteq \Gamma^{(k)}_{G\setminus e}$. Moreover every $\gamma \in \Gamma^{(k)}_{G\setminus e} \setminus \Gamma^{(k)}_G$ has $\gamma(h(e)) = \gamma(t(e))$ and thus induces a valid coloring γ' of G/e via $\gamma' : V' \to \mathbb{N}$ with $\gamma'([v]_\sim) = \gamma(v)$ for all v. Conversely, every $\gamma' \in \Gamma^{(k)}_{G/e}$ induces a unique $\gamma \in \Gamma^{(k)}_{G\setminus e}$. Therefore $\Gamma^{(k)}_{G\setminus e} \setminus \Gamma^{(k)}_G$ and $\Gamma^{(k)}_{G/e}$ have the same cardinality, and thus $|\Gamma^{(k)}_{G\setminus e}| = |\Gamma^{(k)}_G| + |\Gamma^{(k)}_{G/e}|$. The claim follows. □

Corollary-Definition 1.2.5. *For every given finite graph* G *there is a polynomial* $\chi_G \in \mathbb{Z}[x]$ *such that* $\chi_G(k) = \chi^{\mathbb{N}}_G(k)$ *for all* $k \in \mathbb{N}$. *This* χ_G *is called the* chromatic polynomial *of* G.

We are going to give what is called a "combinatorial interpretation" of the coefficients of the chromatic polynomial of a graph. This means that, although we know that we can write

$$\chi_G(t) = \sum_{i=0}^{n} a_i t^i \tag{1.1}$$

for some integers a_i, we would like to be able to determine such coefficients. We present one of the earliest results in this vein, obtained by Hassler Whitney in 1934 [8]. We will follow Whitney's original argument, using an elementary but extremely versatile tool in enumerative combinatorics.

Principle of inclusion-exclusion. Let $(X_i)_{i\in I}$ be a family of (not necessarily distinct) subsets of a finite set X. Then

$$\left|X \setminus \left(\bigcup_{i\in I} X_i\right)\right| = |X| - \sum_{i\in I} |X_i| + \sum_{\substack{i,j\in I\\ i\neq j}} |X_i \cap X_j| - \ldots\ldots + (-1)^{|I|}\left|\bigcap_{i\in I} X_i\right|.$$

Proposition 1.2.6. *Let* G *be a graph with* n *vertices. Then the chromatic polynomial satisfies*

$$\chi_G(t) = \sum_{A\subseteq E} (-1)^{|A|} t^{c(G[A])}. \tag{1.2}$$

In particular, if we let $g_{p,q}$ *denote the number of* $A \subseteq E$ *such that* $|A| = p$ *and* $c(G[A]) = q$, *then the coefficients of the expansion in Equation* (1.1) *can be written as* $a_i = \sum_{j=0}^{n} (-1)^j g_{i,j}$ *for all* i.

Proof. Fix a positive integer k and let us call k-*precoloring* of G any function $\tilde{\gamma} : V \to [k]$, i.e., an *arbitrary* assignment of one of the k colors to the vertices of G. There are $k^{|V|}$ precolorings of G. Such a precoloring will sometimes not be a valid coloring of G, as there might be some "bad" edges of G joining

two vertices of the same color. Given $A \subseteq E$, we let $\mathrm{Pre}(A)$ be the set of all precolorings $\widetilde{\gamma}$ whose set of bad edges is A.
Formally, we set $\mathrm{Pre}(A) := \{\widetilde{\gamma} : V \to \mathbb{N}_{>0} \mid \widetilde{\gamma}(h(e)) = \widetilde{\gamma}(t(e)) \forall e \in A\}$. Notice that

(1) $\mathrm{Pre}(A_1) \cap \mathrm{Pre}(A_2) = \mathrm{Pre}(A_1 \cup A_2)$ for all $A_1, A_2 \subseteq E$,

(2) $|\mathrm{Pre}(A)| = k^{|c(G[A])|}$.

The idea is now to count all precolorings of G that do not have any bad edges, i.e., to enumerate the set

$$\Gamma_G(k) = k^V \setminus \bigcup_{e \in E} \mathrm{Pre}(e)$$

We can apply the principle of inclusion-exclusion to our situation, and write

$$\begin{aligned} \chi_G(k) = |\mathrm{Pre}(\emptyset)| &= k^{|V|} - \sum_{\emptyset \subsetneq A \subseteq E} (-1)^{|A|} |\mathrm{Pre}(A)| \\ &= \sum_{A \subseteq E} (-1)^{|A|} k^{c(A)}, \end{aligned} \tag{1.3}$$

where in the last equality we used that $|V| = c(G[\emptyset])$, hence $k^{|V|} = (-1)^{\emptyset} k^{c(G[\emptyset])}$. The claim follows. □

Definition 1.2.7. Let G be a finite graph and fix any total order $<$ on its set E of edges. A *broken circuit* of G is any set of edges that is obtained by deleting the maximal element of any circuit. More precisely, a broken circuit is any set of the form $C \setminus \max_< C$ where C is a circuit of G.

A *no-broken-circuit set* (or just an "nbc-set") of G is any $A \subseteq E$ that does not contain any broken circuit. Let $\mathrm{nbc}(G)$ denote the family of all nbc-sets of G and for every i let $\mathrm{nbc}_i(G)$ denote the number of nbc-sets of cardinality i.

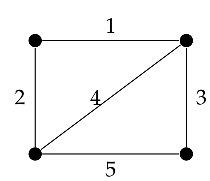

Figure 5: The graph for Example 1.2.8

Example 1.2.8. Consider the graph in Figure 5. Its chromatic polynomial has been computed in Figure 4 and can be expanded to

$$\chi_G(t) = t(t-1)(t-2)^2 = t^4 - 5t^3 + 8t^2 - 4t$$

We identify the set of edges with $\{1, 2, \ldots, 5\}$ according to the numbering displayed in Figure 5. We consider on this set the natural total order of integer numbers.

This graph has three circuits, with edge-sets 124, 345 and 1235. Therefore, there are three broken circuits: 12, 34, 123. The no-broken-circuit sets are then

Size 0:	$\emptyset$	$\mathrm{nbc}_0(G) = 1$
Size 1:	1, 2, 3, 4, 5	$\mathrm{nbc}_1(G) = 5$
Size 2:	13, 14, 15, 23, 24, 25, 35 45	$\mathrm{nbc}_2(G) = 8$
Size 3:	135, 145, 235, 245	$\mathrm{nbc}_3(G) = 4$
Size 4:	none	$\mathrm{nbc}_4(G) = 0$

Theorem 1.2.9. *Let G be a graph on n vertices. Then the coefficients of the chromatic polynomial of G as written in Equation* (1.1) *satisfy*

$$a_{n-i} = (-1)^i \,\mathrm{nbc}_i(G)$$

for all i.

Proof. Recall the total ordering $<$ on E. We extend this to a total ordering on the set of all broken circuits, with the property that $B < B'$ implies $\max_<(B) \leqslant \max_<(B')$, for any two broken circuits B, B'. ([1]) We can now enumerate all broken circuits in this ordering as $B_1, B_2, \ldots, B_s$ and define a partition of the set of subsets of the edge set E as

$$2^E = \mathscr{E}_1 \sqcup \ldots \sqcup \mathscr{E}_{s+1}, \text{ with}$$
$$\mathscr{E}_1 := \{A \subseteq E \mid B_1 \subseteq A\} \text{ and } \mathscr{E}_i := \{A \subseteq E \mid B_i \subseteq A\} \setminus \mathscr{E}_{i-1} \text{ for } i > 1,$$

so the elements of $\mathscr{E}_i$ are all subsets that contain B_i but none of the B_j for $j < i$. Let us now consider the summands in Equation (1.3) subdivided according to the $\mathscr{E}_i$s:

$$\chi_G(k) = \left(\sum_{A \in \mathscr{E}_1} (-1)^{|A|} k^{c(G[A])}\right) + \ldots + \left(\sum_{A \in \mathscr{E}_{s+1}} (-1)^{|A|} k^{c(G[A])}\right)$$

For every $i = 1, \ldots, s$ now let e_i be the edge that has been removed in order to form the broken circuit B_i. The ordering that we have chosen is such that

[1] There are several such total orderings - for instance, we can consider the *lexicographic ordering*, which is defined by setting $B < B'$ if and only if $\max_<(B \setminus B' \cup B' \setminus B) \in B'$. (The name derives from the following alternative definition of this total order: for every broken circuit B consider the ordered list $\lambda_1(B) > \lambda_2(B) > \ldots$ of its elements in decreasing order, and set $B < B'$ if $\lambda_i(B) < \lambda_i(B')$ for the smallest i where the lists differ.)

$e_i \notin B_j$ for all $j < i$ (in fact, $e_i > \max B_i \geqslant \max B_j$). Therefore, if we consider the bipartition

$$\mathscr{E}_i = \underbrace{\{A \in \mathscr{E}_i \mid e_i \in A\}}_{=:\mathscr{E}_i^+} \sqcup \underbrace{\{A \in \mathscr{E}_i \mid e_i \notin A\}}_{=:\mathscr{E}_i^-},$$

the function $\epsilon : \mathscr{E}_i^+ \to \mathscr{E}_i^-$, $\epsilon(A) := A \setminus \{e_i\}$, is a bijection.

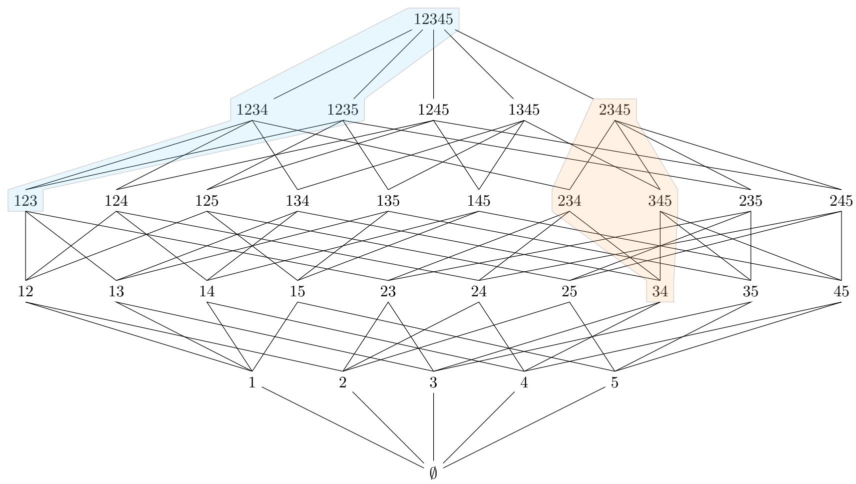

Figure 6: In the situation of Example 1.2.8 the construction of the proof of Theorem 1.2.9 would consider the broken circuits $B_1 = 12$, $B_2 = 123$, $B_3 = 34$, leading to the following families of sets: $\mathscr{E}_1$ as shaded blue in the picture, $\mathscr{E}_2 = \emptyset$, $\mathscr{E}_3$ as shaded orange in the picture, $\mathscr{E}_4 = \mathrm{nbc}(G)$.

Notice that, since the ends of e_i are already connected by the path B_i in each $A \in \mathscr{E}_i^-$, by applying ϵ the number of connected components does not change: $c(G[\epsilon(A)]) = c(G[A])$, but the cardinality decreases by one. Therefore, the contributions of the summands A and $\epsilon(A)$ cancel out in the sum:

$$\sum_{A \in \mathscr{E}_i} (-1)^{|A|} k^{c(G[A])} = \sum_{A \in \mathscr{E}_i^+} \left((-1)^{|A|} k^{c(G[A])} + (-1)^{|\epsilon(A)|} k^{c(G[\epsilon(A)])} \right) = 0.$$

Therefore, only the sum associated to $\mathscr{E}_{s+1}$ contributes to the expression of $\chi_G(k)$. Now notice that in fact $\mathscr{E}_{s+1} = \mathrm{nbc}(G)$. Moreover, no $A \in \mathscr{E}_{s+1}$ contains any circuit of G, and thus $c(G[A]) = |V| - |A|$ for such A. We can now rewrite

$$\chi_G(k) = \sum_{A \in \mathscr{E}_{s+1}} (-1)^{|A|} k^{c(G[A])} = \sum_{i=0}^{n} (-1)^i \, \mathrm{nbc}_i(G) k^{n-i}$$

as desired. □

1.3 Flows

Definition 1.3.1. Let $G = (V, E, h, t)$ be a graph. An *integer flow* on G is any function $f : E \to \mathbb{Z}$ such that for every $v \in V$ we have

$$\sum_{e \in h^{-1}(v)} f(e) - \sum_{e \in t^{-1}(v)} f(e) = 0. \tag{1.4}$$

Let $k \in \mathbb{N} \setminus \{0\}$. We call f an *integer* k*-flow* if $|f(e)| < k$ for all $e \in E$.

A *modular* k*-flow* is any $f : E \to \mathbb{Z}_k$ satisfying Equation (1.4) understood as a congruence modulo k (i.e., as an equation in the Abelian group $\mathbb{Z}_k$).

Any flow f is called *nowhere zero* if $0 \notin f(E)$. The set of all nowhere zero integer k-flows, resp. nowhere zero modular k-flows, is denoted by $\Phi_G^{(k)}$, resp. $\overline{\Phi}_G^{(k)}$.

Remark 1.3.2.

- The definition of a flow on a graph can be given with values in an arbitrary Abelian group. Relevant examples include real-valued flows, of interest in optimization of networks, or complex-valued flows, arising from the analysis of electrical networks. In fact, Equation (1.4) is easily seen to correspond to Kirchoff's conservation condition for flows in networks.

Definition 1.3.3. Let G be a finite graph. The integer flow function and the modular flow function of G are defined as

$$\varphi_G^{\mathbb{N}} : \mathbb{N} \to \mathbb{N}, \quad k \mapsto |\Phi_G^{(k)}|, \qquad \overline{\varphi}_G^{\mathbb{N}} : \mathbb{N} \to \mathbb{N}, \quad k \mapsto |\overline{\Phi}_G^{(k)}|,$$

and they "count" the number of nowhere-zero integer k-flows, resp. nowhere zero modular k-flows on G.

Example 1.3.4.

(1) If G has no edges, clearly $\overline{\varphi}_G^{\mathbb{N}}(k) = \varphi_G^{\mathbb{N}}(k) = 1$ for every k

(2) If e is a loop of G, then $\overline{\varphi}_G^{\mathbb{N}}(k) = (k-1)\overline{\varphi}_{G\setminus e}^{\mathbb{N}}(k)$.

Lemma 1.3.5. *If* G *has an isthmus, then* $\varphi_G^{\mathbb{N}}(k) = \overline{\varphi}_G^{\mathbb{N}}(k) = 0$ *for all* k.

Proof. Let e be an isthmus of G and let f be a flow on G. Let K be the connected component of G containing e, and let S be the vertex set of one of the two connected components of K that arise deleting e. Then, up to sign, $f(e)$ is the total flow $f(S, V \setminus S)$ and thus, by Exercise I.7, equals zero. □

Theorem 1.3.6. *Let* G *be a finite graph and let* e *be an edge of* G*. If* e *is an isthmus, then* $\overline{\varphi}_G^{\mathbb{N}} \equiv 0$*. If* e *is not a loop nor an isthmus of* G*,*

$$\overline{\varphi}_G^{\mathbb{N}}(k) = \overline{\varphi}_{G/e}^{\mathbb{N}}(k) - \overline{\varphi}_{G\setminus e}^{\mathbb{N}}(k) \quad \textit{for every } k. \tag{1.5}$$

Together with Lemma 1.3.5 and Example 1.3.6.(1), this relation determines the function $\overline{\varphi}_G^{\mathbb{N}}$ uniquely.

Proof. Let v, w be the ends of e, and set

$$\Sigma_v := \sum_{\substack{e \in E \setminus e \\ h(e)=v}} f(e) - \sum_{\substack{e \in E \setminus e \\ t(e)=v}} f(e), \quad \Sigma_w := \sum_{\substack{e \in E \setminus e \\ h(e)=w}} f(e) - \sum_{\substack{e \in E \setminus e \\ t(e)=w}} f(e).$$

Now any $f \in \overline{\Phi}^{(k)}_{G\setminus e}$, viewed as a function on the edges of G/e, satisfies trivially Equation (1.4) at every vertex of G/e other than the vertex $[v]_\sim$, i.e., the one formed merging v and w. At this vertex, Equation (1.4) in G/e is $\Sigma_v + \Sigma_w = 0$ and this is trivially satisfied since f, being a flow on $G \setminus e$, satisfies both $\Sigma_v = 0$ and $\Sigma_w = 0$. Thus, $\overline{\Phi}^{(k)}_{G\setminus e} \subseteq \overline{\Phi}^{(k)}_{G/e}$.

On the other hand, an $f \in \overline{\Phi}^{(k)}_{G/e}$ that is not in $\overline{\Phi}^{(k)}_{G\setminus e}$ is one for which $\Sigma_v = -\Sigma_w \neq 0$. This means that it extends uniquely to a nowhere-zero flow $f' \in \overline{\Phi}^{(k)}_G$ by setting, for every $g \in E$,

$$f'(g) := \begin{cases} f(g) & g \neq e \\ \Sigma_{t(e)} & g = e \end{cases}$$

It remains to check that, in fact, $\Sigma_v = -\Sigma_w \neq 0$ must hold also for every nowhere-zero flow on G, so that indeed there is a bijection between $\overline{\Phi}^{(k)}_{G/e} \setminus \overline{\Phi}^{(k)}_{G\setminus e}$ and $\overline{\Phi}^{(k)}_G$. The claim follows now in view of the already proved containment $\overline{\Phi}^{(k)}_{G\setminus e} \subseteq \overline{\Phi}^{(k)}_{G/e}$. □

Corollary-Definition 1.3.7. *For every given finite graph* G *there is a polynomial* $\overline{\varphi}_G \in \mathbb{Z}[x]$ *such that* $\overline{\varphi}_G(k) = \overline{\varphi}_G^{\mathbb{N}}(k)$ *for all* $k \in \mathbb{N}$. *This* $\overline{\varphi}_G$ *is called the* flow polynomial *of* G.

Example 1.3.8. We can compute the flow polynomial for the graphs in Figures 2 and 4 and compare them with the associated chromatic polynomial.

G	Figure 2a	Figure 2b	Figure 4a	Figure 4b
$\chi_G(t)$	$t(t-1)(t-2)$	$t(t-1)$	$t(t-1)(t-2)^2$	$t(t-1)(t-2)$
$\overline{\varphi}_G(t)$	$(t-1)$	$(t-1)(t-2)$	$(t-1)(t-2)$	$(k-1)(k-2)^2$

We have seen that for some graphs, for example K_4, the chromatic and flow polynomials are related up to a multiplication by the variable. In the table above we see two pairs graphs where the chromatic polynomial of the one is the flow polynomial of the other (again, up to a multiplication of the variable), and vice-versa. Those graphs are related by what is called “planar duality”.

Digression: planar duality of graphs.

Let $G = (V, E, t, s)$ be a graph. A *planar drawing* of G is a set of points and paths in the plane satisfying some properties. More precisely, a drawing consists of

- one point $p_v \in \mathbb{R}^2$ for each $v \in V$, so that the function $v \mapsto p_v$ is injective and p_V is discrete in $\mathbb{R}^2$.
- a simple, regular path ω_e for every $e \in E$, i.e., an injective and continuous function $\omega_e : [0,1] \to \mathbb{R}^2$ with $\omega_e(0) = p_{t(e)}$ and $\omega_e(1) = p_{h(e)}$, and for which $\nabla\omega_e(t) \neq 0$ for all $t \in]0,1[$.

We require that the paths representing edges only intersect at points labeled by the vertices incident to the edge under consideration. Precisely: for every $e, e' \in E$, we require that $\omega_e(t) = \omega_e(t')$ implies $\{t, t'\} \subseteq \{0, 1\}$.

This definition entails that if C is the set of edges of a circuit of G, the set $\bigcup_{e \in C} \omega_e([0,1])$ is a Jordan curve. Thus, by the Jordan curve theorem every circuit of G determines a bipartition of $\mathbb{R}^2$ in two open and connected regions (one bounded and one unbounded). The (open, connected) parts of the partition of $\mathbb{R}^2$ that refines all bipartitions obtained by circuits of G are called the *regions* determined by the given drawing of G. They are all bounded and contractible, with exception of one unbounded (and in general not-contractible) region. The boundary of every region is a circuit of G drawn in the plane.

We can define now a planar dual G^* to G as follows. The vertices of G^* are the regions of the drawing, and the edge-set can be taken to be the same as the edge-set of G, but for clarity we will call it E^*, with a fixed bijection $e \mapsto e^*$ between E and E^*. The ends of any e^* are determined as follows: look at the drawing ω_e of the edge e and notice that $\omega_e([0,1])$ is contained in the (topological) boundary of at most two regions, say R_1, R_2 (where we will allow $R_1 = R_2$). As we go along the boundary of R_1 and R_2 in positive direction (counterclockwise), we'll meet in one case (say R_1) first $t(e)$ and then $h(e)$, in the other (say R_2) the reverse will happen. We will then say that $t^*(e^*) = R_1$, $h^*(e^*) = R_2$.

Alternatively (and most commonly), one obtains a drawing of G^* by choosing a point (representing a vertex of G^*) in each region of the given drawing of G and drawing a simple regular path (representing an edge e^* of G^*) across each edge e of G, whose endpoints are the vertices of the region(s) in whose boundary $\omega_e([0,1])$ lays, and that crosses the path representing e at a unique crossing point $p(e)$. The ends $t(e^*)$ and $h(e^*)$ are assigned so that the pair $(\nabla\omega_e, \nabla\omega_{e^*})$ taken at $p(e)$ is a positively oriented basis of the plane. (intuitively: "e^* crosses e from left to right", see Figure 7).

Different drawings of the same abstract graph G can give rise to non-isomorphic duals. However, it is always true that $(G^*)^*$ is isomorphic to G.

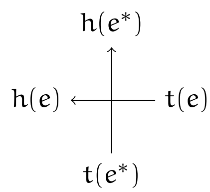

Figure 7: Illustration of the "orientation" of primal-dual edges

Proposition 1.3.9. *Let* G *be a finite, connected planar graph and let* G^* *be its planar dual. For every* $k \in \mathbb{N} \setminus \{0\}$*, there is a* k*-to-one surjective map between the set* $\Gamma_G^{(k)}$ *of* k*-colorings of* G *and the set* $\overline{\Phi}_{G^*}^{(k)}$ *of nowhere-zero* k*-flows on* G^**.*

Corollary 1.3.10. *Let* G *be a finite, connected planar graph. Then* $\chi_G(t) = t\overline{\varphi}_{G^*}(t)$*.*

Proof of Proposition 1.3.9. Let G be any given finite planar graph and consider its planar dual G^*. In particular, every edge e of G is paired with an edge e^* of G^*, i.e., the one that crosses e (at a unique crossing point $p(e)$) in the given planar drawing of G.

Now, given a k-coloring γ of G define a function $f^* : E^* \to \mathbb{Z}_k$ as follows:

$$f_\gamma(e^*) := \gamma(t(e)) - \gamma(h(e)),$$

where on the r.h.s. we take the congruence class modulo k. It is easy to see that γ being a proper coloring implies that f_γ is nowhere zero. In order to check the condition for f_γ to be a flow, consider any vertex v^* of G^*. This vertex corresponds to a region of the plane that is bounded by a circuit of G. Call $C : v_1, e_1, v_2, \ldots$ this circuit, with vertices and edges listed in counterclockwise direction. Then every edge of G^* that insists on v^* is of the form $(e_i)^*$ for some e_i in C, and by construction of the planar dual we have that $v^* = t(e_i^*)$ if and only if $v_i = t(e_i)$ (and thus $v_{i+1} = h(e_i)$), see Figure 8.

Now we can compute Equation (1.4)

$$\sum_{v^*=t(e^*)} f_\gamma(e^*) - \sum_{v^*=h(e^*)} f_\gamma(e^*) =$$
$$\sum_{v_i=t(e_i)} \gamma(t(e_i)) - \gamma(h(e_i)) + \sum_{v_i=h(e_i)} \gamma(h(e_i)) - \gamma(t(e_i)) =$$
$$\sum_{v_i=t(e_i)} \gamma(v_i) - \gamma(v_{i+1}) + \sum_{v_i=h(e_i)} \gamma(v_i) - \gamma(v_{i+1}) = \sum_{v_i} \gamma(v_i) - \gamma(v_i) = 0$$

Now, fix a vertex $v_0 \in V$. Let Γ be the set of all k-colorings of G and let Φ be the set of all nowhere-zero k-flows on G^*. We already know that the function

$$\Gamma_G^{(k)} \to \overline{\Phi}_G^{(k)} \times \{1, \ldots, k\}, \quad \gamma \mapsto (f_\gamma, \gamma(v_0))$$

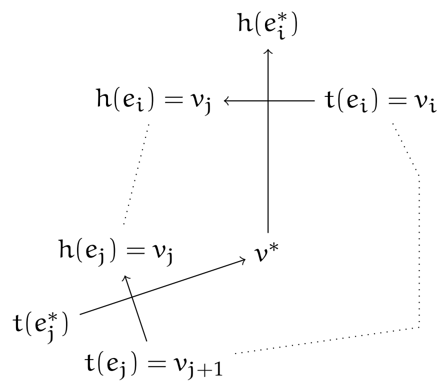

Figure 8: Here $v^* = t(e_j^*) = h(e_i^*)$

is well-defined. We are left with proving that it is bijective, which we will do by providing an inverse. Let f be a nowhere-zero k-flow on G^* and let $i_0 \in [k]$. For every vertex v of G choose a path $v_0, e_1, v_1, \ldots, v_l = v$. Define

$$\gamma_f(v) \in \{1, \ldots k\}, \quad \gamma_f(v) \equiv i_0 + \sum_{h(e_i)=v_i} f((e_i)^*) - \sum_{t(e_i)=v_i} f((e_i)^*) \mod k$$

Notice that $\gamma_f(v)$ does not depend on the choice of paths (we prove this as Exercise I.8) and is thus well-defined. Now notice that for any edge e we have $\gamma_f(h(e)) - \gamma_f(t(e)) = f(e^*)$ and thus the fact that f is nowhere-zero implies that γ_f is a proper k-coloring.

□

Since the double dual of a planar graph is the graph itself, Corollary 1.3.10 can be rewritten as saying that $\overline{\varphi}_G(t) = t^{-1}\chi_{G^*}(t)$ for any connected planar graph G. Thus Proposition 1.2.6 can give us an explicit expression for the flow polynomial of a connected planar graph, provided we can interpret $c(G^*[A])$ in terms of G.

Definition 1.3.11. Let G be any finite graph. Define the *cyclomatic number* of G to be

$$\beta_1(G) := |E| - |V| + c(G).$$

As the name suggests, this is the first Betti number of the graph G viewed as a 1-dimensional simplicial complex. If G is planar, this is one less than the number of faces.

Lemma 1.3.12. *Let G be any finite graph, e an edge of G. Then*

$$\beta_1(G \setminus e) = \begin{cases} \beta_1(G) & \text{if } e \text{ is an isthmus,} \\ \beta_1(G) - 1 & \text{otherwise.} \end{cases} \tag{1.6}$$

$$\beta_1(G/e) = \begin{cases} \beta_1(G) - 1 & \textit{if } e \textit{ is a loop,} \\ \beta_1(G) & \textit{otherwise.} \end{cases} \tag{1.7}$$

Theorem 1.3.13. *Let* G *be a finite graph. The modular flow polynomial of* G *can be written as*

$$\overline{\varphi}_G(t) = (-1)^{|E|} \sum_{A \subseteq E} (-1)^{|A|} t^{\beta_1(G[A])}$$

Proof. Let $P_G(t)$ denote the sum on the r.h.s. in the claim. We have to show that it equals $\overline{\varphi}_G(t)$. We argue using Theorem 1.3.6, which determines the flow polynomial fully.

- If G has no edges, then the sum that defines $P_G(t)$ has only one summand (for $A = \emptyset$) and the cyclomatic number is 0. Thus, $P_G(t) = (-1)^0(-1)^0 t^0 = 1$.

- Let e be an isthmus of G. With Lemma 1.3.12 we compute

$$\begin{aligned} P_G(t) &= (-1)^{|E|} \sum_{A \subseteq E \setminus \{e\}} \left((-1)^{|A|} t^{\beta_1(G[A])} + (-1)^{|A \cup \{e\}|} t^{\beta_1(G[A \cup \{e\}])} \right) \\ &= (-1)^{|E|} \sum_{A \subseteq E \setminus \{e\}} \left((-1)^{|A|} t^{\beta_1(G[A])} - (-1)^{|A|} t^{\beta_1(G[A])} \right) = 0 \end{aligned}$$

- Let e be a loop of G. Again, we compute with Lemma 1.3.12:

$$\begin{aligned} P_G(t) &= (-1)^{|E|} \sum_{A \subseteq E \setminus \{e\}} \left((-1)^{|A|} t^{\beta_1(G[A])} + (-1)^{|A \cup \{e\}|} t^{\beta_1(G[A \cup \{e\}])} \right) \\ &= (-1)^{|E|} \sum_{A \subseteq E \setminus \{e\}} (1-t)(-1)^{|A|} t^{\beta_1(G[A])} \\ &= (t-1)(-1)^{|E \setminus \{e\}|} \sum_{A \subseteq E \setminus \{e\}} (-1)^{|A|} t^{\beta_1(G[A])} = (t-1) P_{G \setminus e}(t) \end{aligned}$$

- Now suppose that e is no loop. Then, contracting e does not change the cyclomatic number. We compute

$$\begin{aligned} P_G(t) &= (-1)^{|E|} \left(\sum_{e \in A \subseteq E} (-1)^{|A|} t^{\beta_1(G[A]/e)} + \sum_{A \subseteq E \setminus \{e\}} (-1)^{|A|} t^{\beta_1(G[A])} \right) \\ &= (-1)^{|E|-1} \left(\sum_{A \subseteq E \setminus \{e\}} (-1)^{|A|} t^{\beta_1(G[A]/e)} - \sum_{A \subseteq E \setminus \{e\}} (-1)^{|A|} t^{\beta_1(G[A])} \right) \\ &= P_{G/e}(t) - P_{G \setminus e}(t). \end{aligned}$$

This completes the proof. □

Exercises I

I.1 Show that a finite graph G with n vertices and m edges is cycle-free if and only if $c(G) = n - m$.

I.2 The *degree* of a vertex v of a finite graph G is the "total number of vertex-edge incidences at v", i.e.,

$$d(v) = |\{e \in E \mid h(e) = v\}| + |\{e \in E \mid t(e) = v\}|.$$

Prove that a finite graph G contains a cycle if and only if it contains a subgraph $G[A]$ whose vertex-degrees are all even.

I.3 Let G be a finite graph and $A, B \subseteq E$ two disjoint subsets of edges of G. Prove that $(G/A) \setminus B = (G \setminus B)/A$.

I.4 Give a formal proof of the principle of inclusion-exclusion.

I.5 Determine the chromatic polynomial of the graph in Figure 9 using deletion-contraction. Check your result using no-broken-circuit sets.

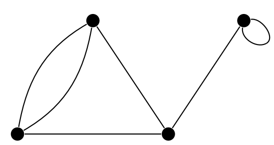

Figure 9

I.6 Determine the modular flow polynomial of the graph in Figure 9 using deletion-contraction. Check your result using no-broken-circuit sets of the planar dual.

I.7 Let $G = (V, E, h, t)$ be a finite graph and let $S, T \subseteq$ be two subsets of its vertex set. Let f be a flow on G with values in an abelian group $\mathbb{A}$, written additively (i.e., we write "+" for the group operation and consider a function $f : E \to \mathbb{A}$ satisfying Equation (1.4) as an identity between group elements with respect to the group operation).

The *total* S-T *flow* is defined as

$$f(S,T) := \sum_{\substack{e \in E \\ t(e) \in S \\ h(e) \in T}} f(e) - \sum_{\substack{e \in E \\ h(e) \in S \\ t(e) \in T}} f(e)$$

Prove that, if $V = S \sqcup T$ is a bipartition of V, then $f(S,T) = 0$.
(Hint: as a first step, try to prove that $f(S,S) = 0$ and $f(S,V) = 0$ for all $S \subseteq V$)

I.8 Prove that the definition of γ_f in the proof of Proposition 1.3.9 does not depend on the choice of paths. (Hint: you may use the previous exercise.)

Chapter 2

The Tutte polynomial

2.1 Definition

From our discussion of colorings and modular flows, we see the importance of two graph statistics: the number of connected components and the cyclomatic number. We can then consider the associated two-variable generating function.

Definition 2.1.1.

$$Q_G(v,w) := \sum_{A\subseteq E} v^{c(G[A])} w^{\beta_1(G[A])} \tag{2.1}$$

Remark 2.1.2. This polynomial, called the "dichromate" of G, generalizes both the chromatic and the flow polynomial of the graph. In fact we have

$$\chi_G(t) = (-1)^{|V|} Q_G(-t,-1), \qquad \overline{\varphi}_G(t) = (-1)^{|E|+|V|} Q_G(-1,-t)$$

Example 2.1.3.

- If G has no edges, then $Q_G(v,w) = v^{|V|}$
- If G has one vertex and a single loop, then $Q_G(v,w) = v(1+w)$
- If G has two vertices and a single isthmus, then $Q_G(v,w) = v(1+v)$

Proposition 2.1.4. *Let G be a finite graph, e an edge of G.*

- *If e is a loop or an isthmus of G, then $Q_G(v,w) = v^{-1} Q_{G[e]}(v,w) Q_{G\setminus e}(v,w)$*
- *If e is neither a loop nor an isthmus of G, then $Q_G(v,w) = Q_{G\setminus e}(v,w) + Q_{G/e}(v,w)$*

Proof. We first expand

$$Q_G(v,w) = \sum_{A\subseteq E\setminus e} v^{c(G[A])} w^{\beta_1(G[A])} + \sum_{e\in A\subseteq E} v^{c(G[A])} w^{\beta_1(G[A])}$$

Now if e is a loop, the second term can be written as

$$w \sum_{e \in A \subseteq E} v^{c(G[A \setminus e])} w^{\beta_1(G[A \setminus e])},$$

whence $Q_G(v, w) = (1 + w) Q_{G \setminus e}(v, w)$ in this case. If e is an isthmus the first term can be written as

$$v \sum_{A \subseteq E \setminus e} v^{c(G[A])} w^{\beta_1(G[A])},$$

and then $Q_G(v, w) = (v + 1) Q_{G \setminus e}(v, w)$.

Otherwise, the second term equals

$$\sum_{e \in A \subseteq E} v^{c(G/e[A \setminus e])} w^{\beta_1(G/e[A \setminus e])}$$

and we conclude $Q_G(v, w) = Q_{G \setminus e}(v, w) + Q_{G/e}(v, w)$.

□

The idea of the dichromate Q_G goes back to W. T. Tutte, but nowadays it is the following variation on it that is commonly called the "Tutte polynomial" of the graph G.

Definition 2.1.5. The *Tutte polynomial* of a finite graph G is

$$\begin{aligned} T_G(x, y) &:= (x-1)^{-c(G)} Q_G(x-1, y-1) \\ &= \sum_{A \subseteq E} (x-1)^{r(G) - r(G[A])} (y-1)^{n(G[A])} \end{aligned}$$

Where we write r and n for the *rank* and *nullity* of a finite graph, defined as

$$r(G) = |V| - c(G), \qquad n(G) = |E| - r(G). \tag{2.2}$$

Remark 2.1.6. Exercise ... shows that the rank is is the cardinality of any maximal cycle-free set of edges of G, and the cyclomatic number equals the number of edges that do not participate in such a cycle-free set (i.e., the nullity of G). In particular, rank and nullity are properties that can be detected without referring to connectivity properties such as isolated vertices etc.

Remark 2.1.7. We can translate Remark 2.1.2 in terms of the Tutte polynomial.

$$\chi_G(t) = (-1)^{r(G)} t^{c(G)} T_G(1-t, 0), \qquad \overline{\varphi}_G(t) = (-1)^{|E| + r(G)} T_G(0, 1-t)$$

2.2 Duality and matroids

Lemma 2.2.1. *Let G be a finite planar graph, and let G^* a planar dual to G. Then,*

$$r(G[A^c]) = r(G) - n(G^*[A^*]).$$

Proof. The equality holds tautologically if A is empty. Now suppose that for some $A \subsetneq E$ the equality holds and choose any $e \in E \setminus A$. Now, $G[(A \cup \{e\})^c] = G[A^c] \setminus e$ and we know that deleting an element drops the rank if and only if that element is an isthmus, i.e., there was a connected component K of $G[A^c]$ that is disconnected when removing e. Now notice that K is the part of G that lies in one of the bounded regions determined by $G^*[A]$, call it R. Now removing e disconnects K exactly when adding e^* to $G^*[A^*]$ subdivides R into two parts, making $n(G^*[A^*])$ increase. This concludes the inductive step and proves the claim. □

Corollary 2.2.2. *If G is a finite planar graph and G^* a dual to G, then*

$$T_{G^*}(x, y) = T_G(y, x).$$

We can now readily re-prove the analog of Corollary 1.3.10 for general planar graphs.

Corollary 2.2.3. *For any finite planar graph G we have*

$$\chi_G(t) = t^{c(G)} \overline{\varphi}_{G^*}(t)$$

These last results depend on the fact that planar graphs have *duals*. The crucial relationship between rank and nullity of primal and dual graphs is expressed in Lemma 2.2.1. From there, recalling the definition of nullity and the fact that the dual to a planar graph has as many edges as the primal graph, we can state that the main "feature" of dual graphs is to provide an interpretation of the quantity

$$r(G[E \setminus A]) - r(G) + |E \setminus A|$$

as the rank function of another graph (i.e., as $r(G^*[A^*])$).

Our next goal is to introduce a class of integer-valued functions that includes rank functions of graphs and that is closed under the above "duality" operation.

Definition 2.2.4. Given any finite set E and any function $r : 2^E \to \mathbb{N}$, let us write

$$r^* : 2^E \to \mathbb{N}, \quad r^*(A) = r(E \setminus A) - r(E) + |A|.$$

Remark 2.2.5. The operation $r \mapsto r^*$ is an involution on the set of all functions $r : 2^E \to \mathbb{N}$ with $r(\emptyset) = 0$. In fact:

$$\begin{aligned}(r^*)^*(A) &= r^*(E \setminus A) - r^*(E) + |A| \\ &= r(A) - r(E) + |A| - r(\emptyset) + r(E) - |E| + |E \setminus A| = r(A) - r(\emptyset) = r(A)\end{aligned}$$

Definition 2.2.6. Let E be a finite set. We call $\mathscr{M}_E$ the set of all functions $r : 2^E \to \mathbb{N}$ satisfying:

(r0) for every $A \subseteq E$: $0 \leqslant r(A) \leqslant |A|$;

(r1) for every $A \subseteq B \subseteq E$: $r(A) \leqslant r(B)$;

(r2) for all $A, B \subseteq E$: $r(A) + r(B) \geqslant r(A \cup B) + r(A \cap B)$.

The elements of $\mathscr{M}_E$ are called "matroid rank functions" on E (or just "matroids on the ground set E").

Remark 2.2.7. Let r be a matroid rank function on the set E and r' a matroid rank function on the set E'. We call these matroids *isomorphic* if there is a bijection $b : E \to E'$ with $r = r' \circ b$.

Theorem 2.2.8. *Let E be a finite set. The set $\mathscr{M}_E$ is closed under duality, i.e., if r is in $\mathscr{M}_E$ so is also r^*.*

Proof. Suppose that r satisfies (r0)– (r2) and consider r^*.

- r^* *satisfies* (r0). Let $A \subseteq E$. By (r2) applied to A and $E \setminus A$ and (r0) we have $r(A) + r(E \setminus A) \geqslant r(E)$, hence with (r1) $r(E \setminus A) - r(E) \geqslant -r(A) \geqslant -|A|$, thus $r^*(A) = r(E \setminus A) - r(E) + |A| \geqslant 0$. On the other hand, (r1) applied to $E \setminus A \subseteq E$ gives $r(E \setminus A) - r(E) \leqslant 0$ and thus $r^*(A) = r(E \setminus A) - r(E) + |A| \leqslant |A|$.
- r^* *satisfies* (r1). Let $A \subseteq B \subseteq E$. Then (r2) applied to $E \setminus B$ and $B \setminus A$, together with (r0), gives $r(E \setminus B) + r(B \setminus A) \geqslant r(E \setminus A)$. With this, we compute

$$\begin{aligned} r^*(B) - r^*(A) &= r(E \setminus B) + |B| - r(E \setminus A) - |A| \\ &\geqslant |B| - |A| - r(B \setminus A) = |B \setminus A| - r(B \setminus A) \geqslant 0 \end{aligned}$$

  where in the last inequality we used (r0) for $B \setminus A$.
- r^* *satisfies* (r2). Let $A, B \subseteq E$. We compute directly

$$\begin{aligned} r^*(A) + r^*(B) &= r(E \setminus A) - r(E) + |A| + r(E \setminus B) - r(E) + |B| \\ &\geqslant r((E \setminus A) \cap (E \setminus B)) + |A| + r((E \setminus A) \cup (E \setminus B)) + |B| - 2r(E) \\ &= r(E \setminus (A \cup B)) - r(E) + |A| + r(E \setminus (A \cap B)) - r(E) + |B| \\ &= r^*(A \cup B) + r^*(A \cap B) + \underbrace{|A| - |A \cup B| - |A \cap B| + |B|}_{=0} \end{aligned}$$

  where the inequality is (r2) for r applied to $E \setminus A$ and $E \setminus B$.

□

Proposition-Definition 2.2.9. *Let G be any finite graph on the edge set E and, for every $A \subseteq E$ write, for brevity, $r_G(A) := r(G[A])$. Then, $r_G \in \mathscr{M}_E$, and we call this the "rank function of the cycle matroid of G".*

Remark 2.2.10. We distinguish two special rank functions. Call r_L the rank function of the graph consisting of a single loop (so $r_L : 2^{\{e\}} \mapsto \mathbb{N}$ with r_L constant equal to 0). Moreover, call r_I the rank function of the graph consisting of a single isthmus (so $r_I : 2^{\{e\}} \mapsto \mathbb{N}$ with $r_I(\{e\}) = 1$).

Remark 2.2.11. Lemma 2.2.1 says that, for every planar graph G, $r_G^* = r_{G^*}$. The class $\mathscr{M}_E$ however contains r_G^* for *all* graphs with edge set E. Thus the notion of duality for matroids extends the planar graph duality to general finite graphs.

We will prove a more general statement than Proposition-Definition 2.2.9, motivated by the following construction.

Definition 2.2.12. Let G be a finite graph on the edge set E with vertex set V. For every $e \in E$ define a vector $w_e \in (\mathbb{F}_2)^V$ as follows. For every $v \in V$ let

$$(w_e)_v := \begin{cases} 1 & \text{if } h(e) \neq t(e) \text{ and } v \in \{h(e), t(e)\} \\ 0 & \text{otherwise} \end{cases}$$

Lemma 2.2.13. *Let* G *be any finite graph with edge set* E *and recall Definition 2.2.12. A given* $A \subseteq E$ *contains a cycle if and only if the set* $\{w_a \mid a \in A\}$ *is linearly dependent over* $\mathbb{F}_2$*.*

Proof. Any $\mathbb{F}_2$-linear dependence in $\{w_a\}_{a \in A}$ is of the form $\sum_{i \in D} w_i = 0$ for some $D \subseteq A$, and this means that every component of the sum $\sum_{i \in D} w_i$ contains an even number of 1s. The parity of the number of 1s in the v-th component is equivalent modulo 2 to the degree of the vertex v in $G[D]$, and thus linear dependencies in $\{w_a\}_{a \in A}$ correspond to subgraphs of $G[A]$ with all-even degrees. Now the claim follows with Exercise I.2, where it is proved that $A \subseteq E$ contains a cycle if and only if $G[A]$ contains a subgraph $G[D]$ with all degrees even. □

Corollary 2.2.14. *Let* G *be any finite graph with edge set* E *and recall Definition 2.2.12. Then, for every* $A \subseteq E$ *we have*

$$r_G(A) = \dim_{\mathbb{F}_2} \operatorname{span}\{w_a \mid a \in A\}.$$

Proof. The corollary follows by noting that $r_G(A)$ is the size of any maximal cycle-free subset of A which, via the previous lemma, is exactly the size of a maximal linear independent subset of $\{w_a\}_{a \in A}$. □

Theorem 2.2.15. *Let* E *be any finite set,* $n \in \mathbb{N}$ *and let* $\mathbb{K}$ *be any field. Consider an* E*-tuple of vectors in the vectorspace* $\mathbb{K}^n$*, i.e., a matrix* $w \in \mathbb{K}^{n \times E}$*. Then the function*

$$r : 2^E \to \mathbb{N}, \qquad A \mapsto \dim_{\mathbb{K}} \operatorname{span}\{w_a \mid a \in A\}$$

is in the set $\mathscr{M}_E$*. We say that the matroid defined by* r *is* representable *over the field* $\mathbb{K}$*.*

Proof. Properties (r0) and (r1) can be proved immediately from the notion of dimension of the span of a finite set of vectors. For (r2) let $A, B \subseteq E$ and recall that, for any two subspaces W_1, W_2 of a given finite-dimensional vectorspace, we have the equality

$$\dim W_1 + \dim W_2 = \dim(W_1 \cap W_2) + \dim(W_1 \cup W_2)$$

Property (r2) now follows by applying this equality to $W_1 = \mathrm{span}\{w_a \mid a \in A\}$, $W_2 = \mathrm{span}\{w_b \mid b \in B\}$ and noticing that

$$\mathrm{span}\{w_e \mid e \in A \cup B\} \subseteq \mathrm{span}\{w_a \mid a \in A\} \cap \mathrm{span}\{w_b \mid b \in B\}.$$

□

We close this section with a discussion of operations on matroids that mimick deletion and contraction of edges in graphs.

Definition 2.2.16. Let $r : 2^E \to \mathbb{N}$ be a matroid rank function, and let $e \in E$. We call e a *loop* of the matroid defined by r if $r(e) = 0$, and we call e an *isthmus* if $r(E \setminus e) = r(E) - 1$.

The rank function of the *restriction* of the matroid to any $A \subseteq E$, written $r_{[A]}$, is just the restriction of r, as a function, to $2^A \subseteq 2^E$ (i.e., $r_{[A]}(X) = r(X)$ for all $X \subseteq A$). The *deletion* of A has rank function $r_{\setminus A} := r_{[E \setminus A]}$ equal to the restriction of r to $E \setminus A$. The rank function of the *contraction* of any $A \subseteq E$ is

$$r_{/A} : 2^{E \setminus A} \to \mathbb{N}, r_{/A}(X) := r(X \cup A) - r(A)$$

Any matroid that is obtained by a sequence of contractions and deletions from a given matroid is called a *minor* of the given matroid.

2.3 Universality

To every matroid rank function $r \in \mathscr{M}_E$ we can associate the *Tutte polynomial*

$$T_r(x, y) := \sum_{A \subseteq E} (x-1)^{r(E)-r(A)} (y-1)^{|A|-r(A)}.$$

Remark 2.3.1. $T_{r_G}(x, y) = T_G(x, y)$ for every graph G, and $T_{r^*}(x, y) = T_r(y, x)$.

Example 2.3.2. For the two rank functions of Remark 2.2.10 we have

$$T_{r_I}(x, y) = x, \qquad T_{r_L}(x, y) = y.$$

Remark 2.3.3. If r and r' are isomorphic matroid rank functions, then $T_r(x, y) = T_{r'}(x, y)$.

The deletion-contraction behaviour of Tutte polynomials of graphs that is inherited from the dichromate generalizes to arbitrary matroids, as follows.

Proposition 2.3.4. *Let* $r \in \mathscr{M}_E$ *be a matroid rank function and let* $e \in E$*. If* e *is either a loop or an isthmus,* $T_r(x,y) = T_{r_{[e]}}(x,y)T_{r_{\setminus e}}(x,y)$*. Otherwise,*

$$T_r(x,y) = T_{r_{\setminus e}}(x,y) + T_{r_{/e}}(x,y).$$

Proof. We leave this proof as an exercise. □

Definition 2.3.5. Let $\mathscr{K}$ be a class of matroid(rank function)s that is closed under minors. A *Tutte-Grothendieck invariant on* $\mathscr{K}$ is any function f defined on $\mathscr{K}$ and with values in a commutative ring R, satisfying

(I) $f(r) = f(r')$ whenever r and r' are isomorphic.

(II) $f(r) = f(r_{\setminus e}) + f(r_{/e})$ if e is not a loop nor an isthmus.

(III) $f(r) = f(r_{[e]})f(r_{\setminus e})$ if e is a loop or an isthmus.

Such an f is called a *generalized* (σ, τ) *Tutte-Grothendieck invariant* for some $\sigma, \tau \in R$ if it satisfies (I) and (III) as well as the following variation of (II):

(II′) $f(r) = \sigma f(r_{\setminus e}) + \tau f(r_{/e})$ if e is not a loop nor an isthmus.

Example 2.3.6. The Tutte polynomial is a Tutte-Grothendieck invariant on any minors-closed class of matroids.

Notation 2.3.7. In the following let $\mathscr{K}$ denote a class of matroids that is obtained from a minor-closed class by removing the empty matroid (i.e., the one with $E = \emptyset$).

Theorem 2.3.8. *Let* R *be a commutative ring and let* $\mathscr{K}$ *be as in Notation 2.3.7. Suppose that* $f : \mathscr{K} \to R$ *is any function satisfying (I) always as well as (II) and (III) whenever* $|E| \geqslant 2$*. Then, for all* $r \in \mathscr{K}$*,*

$$f(r) = T_r(f(r_I), f(r_L)).$$

Proof. For every matroid $r \in \mathscr{K}$, a finite number of applications of (II) and (III) allows to decompose both sides of the claimed equality into the same sum of terms associated to matroids with one element, i.e., with $|E| = 1$. But in this case either $r = r_I$ or $r = r_L$. The claim now follows with Example 2.3.2, where we see that $T_{r_L}(f(r_I), f(r_L)) = f(r_L)$ and $T_{r_I}(f(r_I), f(r_L)) = f(r_I)$. □

A similar argument proves the following universality result for generalized Tutte-Grothendieck invariants.

Theorem 2.3.9. *Let* σ *and* τ *be non-zero elements of a field* $\mathbb{F}$*. Then there is a unique* (σ, τ)*-generalized Tutte-Grothendieck invariant* f *on* $\mathscr{K}$ *with values in* $\mathbb{F}[x, y]$ *and such that* $f(r_I) = x$, $f(r_L) = y$*. This* f *is given explicitly by*

$$f(r) = \sigma^{|E|-r(E)}\tau^{r(E)}T_r\left(\frac{x}{\tau}, \frac{y}{\sigma}\right). \tag{2.3}$$

Proof. One checks explicitly that the right-hand side of Equation (2.3) is indeed a generalized (σ, τ) Tutte-Grothendieck invariant. I.e., for (II′) let $e \in E$ be neither an isthmus nor a loop, so that $r(E) = r_{\setminus e}(E') = r_{/e}(E') + 1$, where we write $E' := E \setminus e$. Then,

$$\tau\left[\sigma^{|E'|-r_{/e}(E')}\tau^{r_{/e}(E')}T_{r_{/e}}(\frac{x}{\tau},\frac{y}{\sigma})\right] + \sigma\left[\sigma^{|E'|-r_{\setminus e}(E')}\tau^{r_{\setminus e}(E')}T_{r_{\setminus e}}(\frac{x}{\tau},\frac{y}{\sigma})\right]$$
$$= \tau\left[\sigma^{|E|-r(E)}\tau^{r(E)-1}T_{r_{/e}}(\frac{x}{\tau},\frac{y}{\sigma})\right] + \sigma\left[\sigma^{|E|-r(E)-1}\tau^{r(E)}T_{r_{\setminus e}}(\frac{x}{\tau},\frac{y}{\sigma})\right]$$
$$= \sigma^{|E|-r(E)}\tau^{r(E)}\left(T_{r_{/e}}(\frac{x}{\tau},\frac{y}{\sigma}) + T_{r_{\setminus e}}(\frac{x}{\tau},\frac{y}{\sigma})\right) = \sigma^{|E|-r(E)}\tau^{r(E)}T_r(\frac{x}{\tau},\frac{y}{\sigma})$$

We leave the verification of (III) as an exercise. The uniqueness part of the claim follows as in the proof of Theorem 2.3.8. □

Example 2.3.10. We have seen that, on the class $\mathscr{K}$ of nonempty graphic matroids, the modular flow function is a $(-1, 1)$-generalized Tutte-Grothendieck invariant, with value 0 on r_I and value $k-1$ on r_L. The universality theorem gives us then directly $\overline{\varphi}(t) = (-1)^{|E|-r(E)}T_r(0, 1-t)$, and thus the expression in Theorem 1.3.13.

Exercises II

II.1 Compute the Tutte polynomial of the graph in Figure 9. Check your computation against the results you obtained in Exercise I.5.

II.2 Show that the graphs in Figure 1 have the same Tutte polynomial but that the associated matroids are not isomorphic (i.e., there is no bijection between the edges of the two that preserves rank).

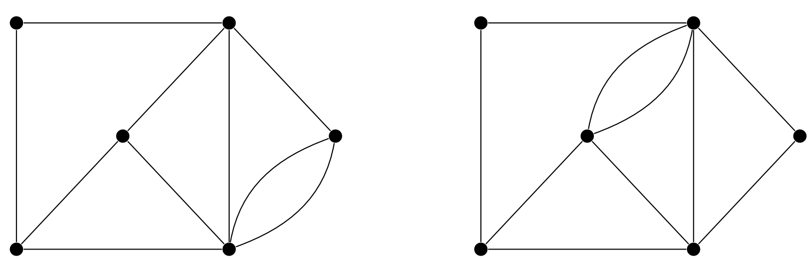

Figure 1

II.3 Let r be the rank function of a matroid on the set E and let $A \subseteq E$. Prove that the dual of the contraction of A equals the deletion of A from the dual (i.e., prove that $(r_{/A})^* = (r^*)_{\setminus A}$).

II.4 Prove Proposition 2.3.4.

II.5 In analogy with the case of graphs, define the *characteristic polynomial* of a given matroid r as $\chi_r(t) := (-1)^{r(E)} T_r(1-t, 0)$.

Chapter 3

Geometric lattices and arrangements of hyperplanes

3.1 Arrangements of hyperplanes

Definition 3.1.1. Let V be vectorspace of finite dimension d. An *arrangement of hyperplanes* in V is a finite set

$$\mathscr{A} := \{H_1, \ldots, H_n\}$$

of codimension 1 linear subspaces of V. The poset of intersections of $\mathscr{A}$ is the set

$$\mathcal{L}(\mathscr{A}) := \left\{ \bigcap_{i \in I} H_i \,\middle|\, I \subseteq [n] \right\}$$

ordered by reverse inclusion: $X \leqslant Y$ if and only if $X \supseteq Y$.

Example 3.1.2. Let $\mathscr{A}$ be the arrengement in $\mathbb{R}^3$ consisting of the four planes

$$\alpha : \{x = 0\}, \quad \beta : \{y = 0\}, \quad \gamma : \{x = y\}, \quad \delta : \{z = 0\},$$

depicted in Figure 1. Then $\mathcal{L}(\mathscr{A})$ is the poset represented on the r.h.s. of Figure 1.

A classical question is the following: given an arrangement of hyperplanes $\mathscr{A}$ in $\mathbb{R}^n$, how many regions are cut out of $\mathbb{R}^n$ by $\mathscr{A}$?

In order to treat this question via the universality of Tutte polynomials, let us set up some theory.

Definition 3.1.3. A partially ordered set P is a *lattice* if, for any two elements $p, q \in P$,

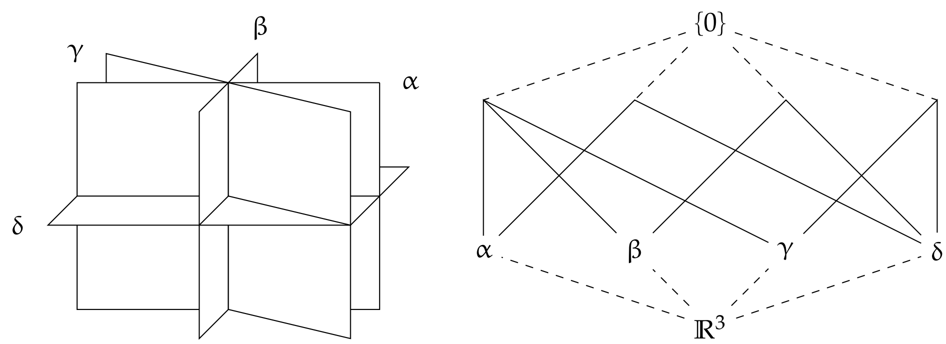

Figure 1

- the subposet $P_{\geqslant p} \cap P_{\geqslant q}$ of all upper bounds to p and q has a unique minimal element - called *join* of p and q and denoted $p \vee q$, and
- the subposet $P_{\leqslant p} \cap P_{\leqslant q}$ of all lower bounds to p and q has a unique maximal element - called *meet* of p and q and denoted $p \wedge q$.

Notice that every finite lattice must have a unique minimal element (denoted by $\hat{0}$) and a unique maximal element (written $\hat{1}$).

Definition 3.1.4. Let P be a poset with a unique minimal element $\widehat{0}$ (we call P "bounded below"). Then the *atoms* of P are the elements of the set

$$A(P) := \{p \in P \mid p \gtrdot \widehat{0}\}.$$

Recall that every finite lattice has a unique minimal element.

Definition 3.1.5. Let L be a finite lattice. We call L *geometric* if, for all $x, y \in L$:

(G) $x \lessdot y$ if and only if there is $p \in A(L)$, $p \nleqslant x$, such that $y = x \vee p$.

Example 3.1.6. Unique least upper bounds exist in $\mathcal{L}(\mathscr{A})$ (for $X, Y \in \mathcal{L}(\mathscr{A})$ take $X \vee Y := X \cap Y$). Moreover, since $\mathcal{L}(\mathscr{A})$ is finite, this implies that also unique greater lower bounds exist (take $X \wedge Y := \bigvee\{Z \in \mathcal{L}(\mathscr{A}) \mid Z \leqslant X, Z \leqslant Y\}$). Thus, $\mathcal{L}(\mathscr{A})$ is a finite lattice.

Now, the atoms of $\mathcal{L}(\mathscr{A})$ are exactly the elements of $\mathscr{A}$, i.e., the hyperplanes. The other nontrivial elements of $\mathcal{L}(\mathscr{A})$ are subspaces of V obtained as intersections of the hyperplanes. Notice here that if W is any linear subspace and H is any hyperplane, the codimension of $H \cap W$ either equals that of W (namely if $H \supseteq W$) or else it surpasses it by one. Therefore, for $W_1, W_2 \in \mathcal{L}(\mathscr{A})$, we have $W_1 \lessdot W_2$ if and only if $W_2 = W_1 \cap H$ for some $H \not\supseteq W_1$ (i.e., $W_2 = W_1 \vee H$ for some $H \in A(\mathcal{L}(\mathscr{A}))$, $H \nleqslant W_1$). In summary, we see that if $\mathscr{A}$ is an arrangement of hyperplanes, then $\mathcal{L}(\mathscr{A})$ is a geometric lattice.

3.2 Matroids from geometric lattices

In what follows we will derive from the definition some properties of a geometric lattice that are "intuitively evident" for intersection posets of hyperplane arrangements. One of these properties is that intersection posets come with a function that assigns to every intersection its codimension as a subspace of V, and this function increases exactly by one along every covering relation. We say that intersection posets are *ranked*. More generally, we have the following definition.

Definition 3.2.1. Let P be a poset. A *rank function* for P is a function $\rho : P \to \mathbb{N}$ such that

(i) $\rho(x) = 0$ if x is a minimal element in P,

(ii) $\rho(x) + 1 = \rho(y)$ if $x \lessdot y$ in P.

Remark 3.2.2. Notice that, if a bounded-below poset admits a rank function, then this function is unique.

Before going forward, let us establish that the *length* of a chain $\omega = \{x_0 < \ldots < x_k\}$ in a partially ordered set P is $\ell(\omega) = |\omega| - 1 = k$. The length of the poset $\ell(P)$ then is the maximum length of any chain in P.

Lemma 3.2.3. *In a geometric lattice any two maximal chains between the same elements have the same length.*

Proof. Let L be a geometric lattice. We prove by induction the following statement (note that in this proof, given $a, b \in L$, an (a, b)-chain is any chain in L of the form $a = x_0 < x_1 < \ldots < x_k = b$).

($*_t$) *For all* $a, b \in L$, *if one maximal* (a, b)*-chain has length* t, *then all of them do.*

The premise of ($*_1$) can only be satisfied if $a \lessdot b$. In this case there is only one maximal (a, b)-chain, hence ($*_1$) holds.

Then let $t \geqslant 2$ and suppose that ($*_r$) holds for all $r < t$. Consider two maximal (a, b)-chains

$$a = c_0 \lessdot c_1 \lessdot \ldots \lessdot c_t = b \qquad a = d_0 \lessdot d_1 \lessdot \ldots \lessdot d_s = b.$$

Now, if $c_1 = d_1$, then by induction hypothesis all maximal (c_1, b)-chains have $t - 1$ elements, hence $s = t$ and we are done.

Suppose then $c_1 \neq d_1$. By property (G) we can find $x, y \in A(L)$ with $c_1 = a \vee x$, $d_1 = a \vee y$. If $x \leqslant d_1$ (resp. $y \leqslant c_1$) we would have $c_1 \leqslant d_1$ (resp. $d_1 \leqslant c_1$), reaching a contradiction; hence, $x \nleqslant d_1$ (resp. $y \nleqslant c_1$). Again by (G), we compute $c_1 \vee d_1 = a \vee x \vee y \gtrdot d_1, c_1$.
Now, by induction hypothesis applied to (c_1, b), every maximal (c_1, b)-chain has length $t-1$, and in particular every maximal $(c_1 \vee d_1, b)$-chain has length $t-2$. In the same way, induction hypothesis applied to (d_1, b) gives that every $(c_1 \vee d_1, b)$-chain has length $s-2$. We conclude $s = t$, and $(*_t)$ holds.

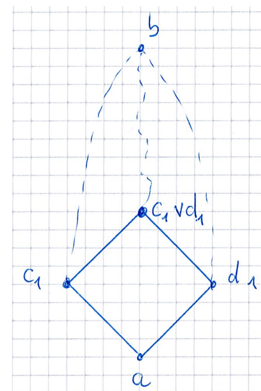

□

Corollary 3.2.4. *Every geometric lattice admits a rank function.*

Proof. Given a geometric lattice L a rank function is given by choosing, for every $x \in L$,

$$\rho(x) := \text{ length of any maximal chain from } \widehat{0} \text{ to } x. \tag{3.1}$$

Lemma Lemma 3.2.3 ensures that this is well-defined, and one readily checks that the conditions of Definition 3.2.1 are satisfied. □

Corollary 3.2.5. *Let* L *be a geometric lattice with rank function* ρ. *For every* $X \subseteq A(L)$ *we have* $\rho(\vee X) \leqslant \#X$.

Proof. First notice that by uniqueness of the rank function we know that ρ can be expressed as in Equation (3.1). Induction on the cardinality of X. If $X = \emptyset$, $\rho(\vee X) = \rho(\widehat{0}) = 0$ and the claim holds.

If $\#X > 0$, choose $x \in X$ and notice that either $\vee(X \setminus \{x\}) = \vee X$ (when $x \leqslant \vee(X \setminus \{x\})$) or, by (G), $\vee(X \setminus \{x\}) \lessdot \vee X$. In any case, a maximal chain from $\widehat{0}$ to $\vee X$ can be obtained by adding at most one new element to a maximal chain from $\widehat{0}$ to $\vee(X \setminus \{x\})$. Therefore, $\rho(\vee X) \leqslant \rho(\vee(X \setminus \{x\})) + 1$ and by induction hypothesis this is at most #X. □

Lemma 3.2.6. *Let* L *be a geometric lattice and* ρ *its*[1] *rank function. Then, for all* $x, y \in L$,

$$\rho(x) + \rho(y) \geqslant \rho(x \wedge y) + \rho(x \vee y).$$

Proof. Consider $z := x \wedge y$ and any saturated chain $z = z_0 \lessdot z_1, \lessdot z_2, \lessdot \cdots \lessdot z_k = y$. Then,

$$k = \rho(y) - \rho(x \wedge y). \tag{3.2}$$

By (G) we can choose atoms $a_1, \dots, a_k$ so that $a_i \leqslant z_i$, $a_i \nleqslant z_{i-1}$ and $z_i = z_{i-1} \vee a_i$ for all $i = 1, \dots, k$.

[1]Unique by Remark 3.2.2

Define now elements $w_0, \ldots, w_k$ by setting $w_0 = x$ and $w_i := w_{i-1} \vee a_i$ for all $i \geqslant 1$. Notice that $w_k = x \vee a_1 \vee \ldots \vee a_k = x \vee z \vee a_1 \vee \ldots \vee a_k = x \vee y$.

Then, by (G) we have either $w_i = w_{i-1}$ or $w_{i-1} \lessdot w_i$ for all i, so that $k \geqslant \rho(w_k) - \rho(w_0) = \rho(x \vee y) - \rho(x)$ and the claim follows by recalling Equation (3.2). □

We have proved the following.

Proposition 3.2.7. *Let E be a finite set and let $\mathcal{L} \subseteq 2^E$ a family of subsets of E, partially ordered by inclusion and such that $E \in \mathcal{L}$. Suppose further that $\mathcal{L}$ is a geometric lattice with rank function ρ.*

Then for every $X \subseteq E$ there is a unique minimal X' in $\mathcal{L}$ such that $X \subseteq X'$, and the extension r of ρ on 2^E given by $r(X) := \rho(X')$ is a matroid rank function.

Proof. Axiom (r1) is trivially satisfied. For Axiom (r0) notice first that ρ is never negative by definition. Moreover, given $X \subseteq E$ we can consider the family $A_1, \ldots, A_k$ of all atoms of $\mathcal{L}$ whose intersection with X is not contained in the minimal element B of $\mathcal{L}$. Then surely $k < |X|$, and

$$X \subseteq B \cup \bigcup_i A_i = \bigcup_i A_i,$$

the second equality because $B \subseteq A_i$ for all i. In particular, we have $X' \leqslant \bigvee_i A_i$ in $\mathcal{L}$, and by Corollary 3.2.5 $\rho(\bigvee_i A_i) \leqslant k$. Thus $r(X) = \rho(X') \leqslant k \leqslant |X|$ as desired.

We now turn to Axiom (r2). First notice that, trivially, $X' \wedge Y' \geqslant (X \cap Y)'$. By definition, $X' \vee Y'$ is the minimal element of $\mathcal{L}$ containing X' and Y', while $(X \cup Y)'$ is the minimal element of $\mathcal{L}$ containing X and Y. Since $X \subseteq X'$ and $Y \subseteq Y'$, we have $X' \vee Y' \geqslant (X \cup Y)'$. With the trivial inequality $X' \vee Y' \leqslant (X \cup Y)'$ we obtain $X' \vee Y' = (X \cup Y)'$.

Now using Lemma 3.2.6 and the monotony of ρ we can write

$$\begin{aligned} r(X) + r(Y) &\overset{\mathrm{df}}{=} \rho(X') + \rho(Y') \\ &\geqslant \rho(X' \wedge Y') + \rho(X' \vee Y') \geqslant \rho((X \cap Y)') + \rho((X \cup Y)') \\ &\overset{\mathrm{df}}{=} r(X \cap Y) + r(X \cup Y) \end{aligned}$$

□

Corollary 3.2.8. *Given any (abstract) geometric semilattice $\mathcal{L}$, we can associate to every $x \in \mathcal{L}$ the set $A(x)$ of all atoms of $\mathcal{L}$ below x. Then, $\mathcal{L}$ is isomorphic to the set $\mathcal{L}' := \{A(x) \mid x \in \mathcal{L}\}$ ordered by inclusion (since $x < y$ if and only if $A(x) \subset A(y)$). The matroid constructed from the proposition, then, has the set $A(\mathcal{L})$ of all atoms as a ground set and rank function given by $r(X) = \rho(\vee X)$ for all $X \subseteq A(\mathcal{L})$. This matroid has no loops, and it is referred to as the "simple matroid associated to $\mathcal{L}$.*

Example 3.2.9. Let us consider the geometric lattice from Figure 1. The set of atoms is $\{\alpha, \beta, \gamma, \delta\}$, and the associated geometric lattice $\mathcal{L}'$ in Corollary 3.2.8 is as follows.

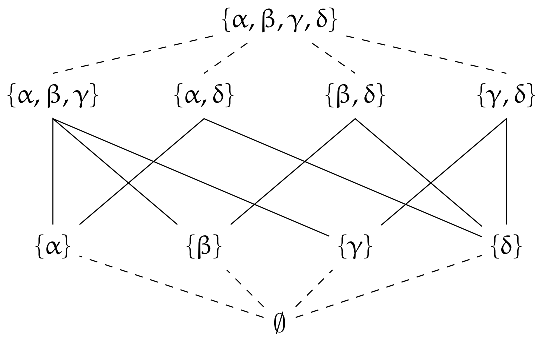

The claim of Corollary 3.2.8 is then that this is the lattice of flats of a matroid on $E = \{\alpha, \beta, \gamma, \delta\}$ with rank function given by $r(A) = |A|$ if $|A| \leqslant 2$, $r(\{\alpha, \beta, \gamma\}) = 2$, and $r(A) = 3$ for all A with $A \neq \{\alpha, \beta, \gamma\}$ and $|A| \geqslant 3$.

3.3 Geometric lattices from matroids

We aim at a "converse" of Proposition 3.2.7, constructing a geometric lattice for every given matroid.

Definition 3.3.1. Let E be a finite set and $r : 2^E \to \mathbb{N}$ a matroid rank function. Define a *closure operator*

$$\mathrm{cl} : 2^E \to 2^E, \quad X \mapsto \{e \in E \mid r(X \cup e) = r(X)\}.$$

Call $X \subseteq E$ *closed* if $X = \mathrm{cl}(X)$, and let $\mathcal{L}_r$ be the poset of all closed sets ordered by inclusion (i.e., for $F, F' \in \mathcal{L}_r$ we have $F \leqslant F'$ if $F \subseteq F'$).

Example 3.3.2.

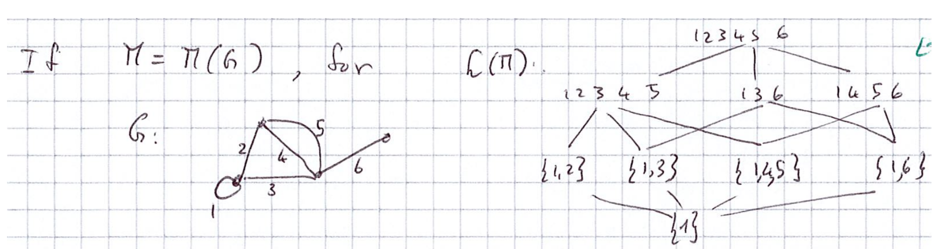

Example 3.3.3. Consider the rank function $r : 2^{[4]} \to \mathbb{N}$ defined by $r(X) = 1$ if $|X| \leqslant 1$ and $r(X) = 2$ otherwise. One can check that this is a matroid rank function – it is called "*uniform of rank* 2 *on* 4 *elements*". The associated poset of flats is depicted below.

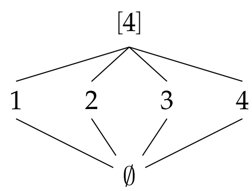

Our next goal is to prove that, in general, $\mathcal{L}_r$ is a geometric lattice.

Remark 3.3.4.

(1) For all $X \subseteq E$ we have obviously $X \subseteq \mathrm{cl}(X)$.

(2) For all $X \subseteq E$, $r(\mathrm{cl}(X)) = r(X)$.
Proof. Let $Y := \mathrm{cl}(X) \setminus X$. We prove that $r(X \cup Y') = r(X)$ for all $Y' \subseteq Y$, by induction on $k := |Y'|$.
The case $k = 1$ holds by definition of "$y \in \mathrm{cl}(X)$".
For $k > 1$, assume that the claim holds for smaller sizes of Y', and choose $y_1, y_2 \in Y'$. Then

$$\begin{aligned} 2r(X) &\overset{\text{I.H.}}{=} r((X \cup Y') \setminus \{y_1\}) + r((X \cup Y') \setminus \{y_2\}) \\ &\overset{r3}{\geqslant} r(X \cup Y') + r(X \cup Y' \setminus \{y_1, y_2\}) \\ &\overset{r2}{\geqslant} r(X) + r(X) \end{aligned}$$

and equality must hold throughout. Thus, $r(X) = r(X \cup Y')$.

(3) For all $X \subseteq E$, $\mathrm{cl}(X) = \max_{\subseteq}\{Y \supseteq X \mid r(Y) = r(X)\}$.
Proof. Let $\Gamma := \{Y \supseteq X \mid r(X) = r(Y)\}$. By (1) and (2) above we have $\mathrm{cl}(X) \in \Gamma$. It is now enough to prove that $Y \subseteq \mathrm{cl}(X)$ for all $Y \in \Gamma$.

To this end, take $Y \in \Gamma$ and, for any $y \in Y$ compute $r(X) \leqslant r(X \cup \{y\}) \leqslant r(Y) = r(X)$ (the last equality because $Y \in \Gamma$), hence $r(X \cup \{y\}) = r(X)$, thus $y \in \mathrm{cl}(X)$. We conclude $Y \subseteq \mathrm{cl}(X)$ as desired.

(4) For all $X \subseteq Y \subseteq E$, $\mathrm{cl}(X) \subseteq \mathrm{cl}(Y)$.
Proof. Clear from (3).

Lemma 3.3.5. *Let r be a matroid rank function. Then, meet and join of every $F_1, F_2 \in \mathcal{L}_r$ exist. In fact,*

(1) $F_1 \vee F_2 = \mathrm{cl}(F_1 \cup F_2)$

(2) $F_1 \wedge F_2 = F_1 \cap F_2$

In particular, $\mathcal{L}_r$ is a lattice.

Proof.

(1) By definition of the ordering, every element of $(\mathcal{L}_r)_{\geqslant F_1} \cap (\mathcal{L}_r)_{\geqslant F_2}$ must contain $F_1 \cup F_2$. But by the definition of the closure operator, $\mathrm{cl}(F_1 \cup F_2)$ is the (unique) smallest closed set containing $F_1 \cup F_2$.

(2) It is enough to prove that $F_1 \cap F_2$ is closed, i.e., that $\mathrm{cl}(F_1 \cap F_2) = F_1 \cap F_2$. We do this next.

- $\supseteq$ This inclusion is clear (see Remark 3.3.4.(1)).
- $\subseteq$ For every $e \in \mathrm{cl}(F_1 \cap F_2)$ we have $r((F_1 \cap F_2) \cup \{e\}) = r(F_1 \cap F_2)$. Hence:

$$\underbrace{r((F_1 \cap F_2) \cup \{e\})}_{=r(F_1 \cap F_2)} + r(F_1) \overset{(r3)}{\geqslant} r(F_1 \cup \{e\}) + r\,(F_1 \cap F_2 \text{ or } (F_1 \cap F_2) \cup \{e\})$$

$$= r(F_1 \cup \{e\}) + r(F_1 \cap F_2)$$

and we conclude $r(F_1) = r(F_1 \cup \{e\})$, hence $e \in F_1$.
Analogously we prove $e \in F_2$, hence $e \in F_1 \cap F_2$.

□

Remark 3.3.6. If $X < Y$ in $\mathcal{L}_r$, then $r(X) < r(Y)$. Otherwise, by r2 we would have $r(X) = r(Y)$ and by definition of closure $Y \subseteq \mathrm{cl}(X) = X$, proving $X = Y$.

Proposition 3.3.7. *For any* $F_1, F_2 \in \mathcal{L}_r$, *(G) holds. I.e.,*

$$F_1 \lessdot F_2 \Leftrightarrow \exists P \in A(\mathcal{L}_r),\ P \nleqslant F_1,\ \text{s.t. } F_2 = F_1 \vee P.$$

Proof.

$\Leftarrow$ Let P be as in the claim. Since P is an atom, $P = \mathrm{cl}(\{e\})$ for some element $e \in E$ and, since $P \nleqslant F_1$, by Remark 3.3.4.(4) it must be $e \in E \setminus F_1$. Now we can write F_2 using Lemma 3.3.5 as

$$F_2 = F_1 \vee P = \mathrm{cl}(F_1 \cup P) = \mathrm{cl}(F_1 \cup \{e\})$$

and we have

$$r(F_1) + r(\{e\}) \geqslant r(\emptyset) + r(F_1 \cup \{e\}).$$

Now, $r(\{e\}) = r(\mathrm{cl}(e))$ and since $P = \mathrm{cl}(e)$ has rank 1, $r(e) = 1$. Thus

$$r(F_2) = r(F \cup \{e\}) \leqslant r(F_1) + 1$$

Moreover, since F_1 is closed and $e \notin F_1$ we have $r(F_1 \cup \{e\}) > r(F_1)$, and we conclude that $r(F_2) = r(F_1) + 1$.

Now by Remark 3.3.6 any $Z \in \mathcal{L}_r$, $F_1 < Z < F_2$, would force $r(F_2) \geqslant r(F_1) + 2$, hence a contradiction. We conclude $F_1 \lessdot F_2$.

- $\Rightarrow$ $F_1 \lessdot F_2$ implies $F_1 \subsetneq F_2$ and so we can choose $e \in F_2 \setminus F_1$. Then $r(\{e\}) = 1$ since otherwise e is in the closure of every flat, in particular we would have $e \in F_1$. It follows that $P := \mathrm{cl}(\{e\})$ is an atom of $\mathcal{L}_r$, and $P \leqslant F_2$ by Remark 3.3.4.(1). Now define

$$F := F_1 \vee P = \mathrm{cl}(F_1 \cup \{e\}).$$

Then the following claim concludes the proof.

Claim. $F_2 = F$.

Proof. We have

$$r(F) \geqslant r(F_1) + 1 = r(F_2). \tag{3.3}$$

The inequality holds since $F \supseteq F_1 \cup \{e\}$, F_1 is closed and $e \notin F_1$, the equality is immediate since $F_1 \lessdot F_2$.

Now since $F_1 \cup \{e\} \subseteq F_2$ Remark 3.3.4.(1) implies $F \subseteq F_2$. Together with Equation (3.3) this shows $F = F_2$.

□

Theorem 3.3.8. *Let* r *be any matroid rank function. Then the poset* $\mathcal{L}_r$ *is a geometric lattice whose rank function* ρ *satisfies* $\rho(F) = r(F)$ *for every* $F \in \mathcal{L}_r$.

Proof. That $\mathcal{L}_r$ is a geometric lattice follows from Lemma 3.3.5 and Proposition 3.3.7. For the claim about rank consider any $F \in \mathcal{L}_r$ and let $\hat{0} \lessdot F_1 \lessdot \ldots \lessdot F_k = F$ be a maximal chain below F. Then, $\rho(F) = k$.

Choose atoms $A_1, \ldots, A_k$ with $F_i = F_{i-1} \vee A_i$ for all i. Since every F_{i-1} is closed and $A_i \not\subseteq F_{i-1}$, we must have

$$r(F_{i-1}) > r(F_{i-1} \cup A_i) = r(F_i) \tag{3.4}$$

(the last equality by Lemma 3.3.5.(1)). On the other hand, (r2) implies

$$r(F_{i-1}) + r(A_i) \geqslant r(\underbrace{F_{i-1} \cap A_i}_{=\hat{0}}) + r(F_{i-1} \cup A_i) = r(F_i) \tag{3.5}$$

and since $r(A_i) = 1$ from Equations (3.4) and (3.5) we conclude $r(F_i) = r(F_{i-1}) + 1$, thus $r(X) = r(F_k) = k = \rho(X)$.

□

3.4 Interlude: Arrangements

At this stage we have two, a priori different, rank functions associated to an arrangement $\mathscr{A} = \{H_1, \dots, H_m\}$ of hyperplanes:

- The rank function r_{lat} of the simple matroid associated to the geometric lattice $\mathcal{L}(\mathscr{A})$ as in Corollary 3.2.8:

$$r_{lat} : 2^{[m]} \to \mathbb{N}, \quad I \mapsto \rho\left(\bigvee_{i \in I} H_i\right)$$

- The rank function r_{dep} obtained via Theorem 2.2.15 from the $[n]$-tuple of vectors $\{n_1, \dots, n_m\}$, where n_i is any choice of normal vector for the hyperplane H_i:

$$r_{dep} : 2^{[m]} \to \mathbb{N}, \quad I \mapsto \dim \operatorname{span}\{v_i \mid i \in I\}.$$

Our next goal is to show that they are the same.

Lemma 3.4.1. *For every intersection* $X \in \mathcal{L}(\mathscr{A})$ *we have* $\rho(X) = \operatorname{codim} X$

Proof. By definition $\rho(X) = k$ means that k is the length of a maximal chain $\widehat{0} \lessdot X_1 \lessdot \cdots \lessdot X_k = X$. Now consider the subspaces X_i. By property (G), every X_i is of the form $X_{i-1} \cap H_i$ for some atom H_i of $\mathcal{L}(\mathscr{A})$ (i.e., hyperplane in $\mathscr{A}$) with $H_i \not\leq X_{i-1}$ (i.e., $H_i \not\supseteq X_{i-1}$). Notice that the latter implies that $X_{i-1} + H_i = \mathbb{R}^d$, the ambient space. Now, elementary linear algebra tells us that

$$\dim(\underbrace{X_{i-1} \cap H_i}_{=X_i}) + \underbrace{\dim(X_{i-1} + H_i)}_{=d} = \dim(X_{i-1}) + \underbrace{\dim(H_i)}_{d-1}$$

and thus $\dim(X_i) = \dim(X_{i-1}) - 1$. Therefore, X has dimension k less than $\widehat{0} = \mathbb{R}^d$, and the proof is complete. □

Proposition 3.4.2.

$$r_{lat} \equiv r_{dep}.$$

Proof. Let $I \subseteq [m]$ and write $X := \bigvee_{i \in I} H_i = \bigcap_{i \in I} H_i$ Then, with Lemma 3.4.1 we know that $r_{lat}(I) = \rho(X) = \operatorname{codim}(X)$. On the other hand, $r_{dep}(I)$ equals the rank of the $d \times |I|$ matrix M whose columns are v_i for i in I. Now, X is the subspace of all points that are orthogonal to each v_i, $i \in I$, and therefore $X = \ker M$. Now, again by elementary linear algebra we know that $\dim \ker M = d - \operatorname{rank} M$. We summarize and conclude

$$r_{dep}(I) = \operatorname{rank} M = d - \dim \ker M = d - \dim X = \operatorname{codim}(X) = r_{lat}(I)$$

□

3.5 Minors

Proposition 3.5.1. *Let* r *be the rank function of a matroid on* E *and let* $e \in E$*. Call* cl, $\mathrm{cl}_{\setminus e}$, $\mathrm{cl}_{/e}$ *the closure operators associated to* r*, its deletion and contraction respectively. Then, for every* $A \subseteq E \setminus e$ *we have*

(1) $\mathrm{cl}_{\setminus e}(A) = \mathrm{cl}(A) \setminus \{e\}$

(2) $\mathrm{cl}_{/e}(A) = \mathrm{cl}(A \cup \{e\}) \setminus \{e\}$

Proof. For item (1) we have that $\mathrm{cl}_{\setminus e}(A)$ is the set of all $x \in E \setminus e$ such that $r_{\setminus e}(A \cup x) = r_{\setminus e}(A)$. The claim follows by noticing that $r_{\setminus e} = r$ on $E \setminus e$.

For item (2) notice that $\mathrm{cl}_{/e}(A)$ is the set of all $x \in E \setminus e$ such that $r_{/e}(A \cup x) = r_{/e}(A)$, i.e., such that $r(A \cup x \cup e) - r(e) = r(A \cup e) - r(e)$. This equality clearly holds if and only if $x \in \mathrm{cl}(A \cup e)$. Since $x \neq e$, claim (2) is proved. □

Theorem 3.5.2. *Let* r *be the rank function of a matroid on the ground set* E *and let* $e \in E$*. Then,* $\mathcal{L}_{r_{\setminus e}}$ *is the sublattice of* $\mathcal{L}_r$ *generated by the joins of all* $x \in E \setminus e$*, and* $\mathcal{L}_{r_{/e}} \simeq (\mathcal{L}_r)_{\geqslant \mathrm{cl}(e)}$.

Proof. Exercise. □

3.6 Dissection theory

Let $\mathscr{A}$ be an arrangement of hyperplanes in $\mathbb{R}^d$ and let $\mathcal{L}(\mathscr{A})$ be the associated geometric lattice of intersections. If we denote by $r_{\mathscr{A}}$ the rank function of the matroid defined by $\mathcal{L}(\mathscr{A})$ (or, equivalently, by the dependency relations of the normal vectors), we obtain a function

$$r : \mathtt{Arr} \to \mathtt{Mat}, \quad \mathscr{A} \mapsto r_{\mathscr{A}}$$

from arrangements to matroids.

Now call $\mathcal{R}(\mathscr{A})$ the set of chambers of $\mathscr{A}$ (so $\mathcal{R}(\mathscr{A}) = \pi_0(\mathbb{R}^d \setminus \cup\mathscr{A})$). Our aim will be to apply matroid theory and the universality of the Tutte polynomial in order to give a formula for computing $|\mathcal{R}(\mathscr{A})|$.

For any given $H \in \mathscr{A}$ consider $\mathscr{A}_H := \mathscr{A} \setminus \{H\}$ and $\mathscr{A}^H := \{H' \cap H \mid H' \in \mathscr{A}_H\}$. We see that the regions of $\mathscr{A}^H$ are “walls” that bisect some of the chambers of $\mathscr{A}_H$. In fact, for every chamber of $\mathscr{A}_H$ that is subdivided in passing to $\mathscr{A}$ there is exactly one chamber of $\mathscr{A}^H$ that performs this "cut". We have shown that

$$|\mathcal{R}(\mathscr{A})| = |\mathcal{R}(\mathscr{A}^H)| + |\mathcal{R}(\mathscr{A}_H)|. \tag{3.6}$$

.

Lemma 3.6.1.

(1) $\mathcal{L}(\mathscr{A}_H)$ *is the sublattice of* $\mathcal{L}(\mathscr{A})$ *generated by the set of atoms* $\mathscr{A} \setminus \{H\}$*.*

(2) $\mathcal{L}(\mathscr{A}^H) \simeq \mathcal{L}(\mathscr{A})_{\geqslant H}$*.*

Proof. Exercise. □

Now notice that $\mathscr{A} = \{H\}$ if and only if $\mathcal{L}_r = \{\widehat{0} < \widehat{1}\}$, and in that case $|\mathcal{R}(\mathscr{A})| = 2$ – in particular, this number only depends on the isomorphism type of $\mathcal{L}_r$. With Equation (3.6) and Lemma 3.6.1 we see that this then is true also for bigger arrangements. Thus, we can write

$$|\mathcal{R}(\mathscr{A})| = f(r_{\mathscr{A}})$$

where $f : \mathtt{Mat} \to \mathbb{Z}$ is a function that is constant on every set of matroids that share the same lattice of flats and that satisfies the deletion-contraction recursion

$$f(r) = f(r_{/e}) + f(r_{\setminus e}) \tag{3.7}$$

Lemma 3.6.2. *Let* $f : \mathtt{Mat} \to \mathbb{Z}$ *be a function that is constant on every set of matroids that share the same lattice of flats and satisfies Equation* (3.7)*. Then,* $f(r_L) = 0$*.*

Proof. Let r be the rank function of the matroid on $E = \{a, b\}$ with $r(\{a\}) = r(\{b\}) = r(\{a, b\}) = 1$. Its flats are $\emptyset$ and E, and so its lattice of flats is isomorphic to that of r_I. Moreover, $r_{/a} = r_L$, and $r_{\setminus a} = r_I$. Now Equation (3.7) gives $f(r_I) = f(r_I) + f(r_L)$, implying $f(r_L) = 0$ as desired. □

Thus, our f satisfies the definition of a Tutte-Grothendieck invariant in $\mathbb{Z}$ (for item (III) in Definition 2.3.5 notice Exercise III.2 and recall that that $f(r_I) = 2$), with $f(r_L) = 0$ and $f(r_I) = 2$. We conclude immediately that

$$|\mathcal{R}(\mathscr{A})| = T_{r_{\mathscr{A}}}(2, 0).$$

3.6.1 An example: graphic arrangements

Let G be a loopless graph with edge-set E and set of vertices V. To G is associated an arrangement $\mathscr{A}_G$ in $\mathbb{R}^V$ with one hyperplane H_e for each $e \in E$, defined via its normal vector n_e that we choose as

$$n_e = (0, \ldots, 0, 1, 0, \ldots, 0, -1, 0, \ldots, 0) \tag{3.8}$$

with nonzero value at entries v and w, where $\{v, w\} = \{h(e), t(e)\}$. Notice that, identifying edges of the graph with hyperplanes of the arrangement, the rank function of the graphic matroid derived from G is the same as that derived from the (lattice of flats of the) arrangement.

Recall further that, associated to a graph, we have a chromatic polynomial

$$\chi_G(t) = (-1)^{r(G)} t^{c(G)} T_G(1-t, 0)$$

where $T_G(x, y)$ is the Tutte polynomial of the (matroid associated to the) graph G.

By our considerations above, we have that the number of regions of $\mathscr{A}_G$ is $|\chi_G(-1)|$. This allows us to easily compute the number of regions of certain classes of such "graphic arrangements". For instance, take $G = K_n$, the complete graph on n vertices. The chromatic polynomial of K_n is easy to compute, and so is the number of regions of the arrangement:

$$\chi_{K_n}(t) = t(t-1)(t-2)\dots(t-n+1), \qquad |\mathcal{R}(\mathscr{A}_{K_n})| = |\chi_{K_n}(-1)| = n!$$

You might recall from Coxeter theory that this is no coincidence – in fact, $\mathscr{A}_{K_n}$ is the set of reflecting hyperplanes of the standard linear representation of the symmetric group (or: Coxeter group of type A_{n-1}), and the chambers are in bijection with the elements of the group... of which there are exactly $n!$.

Question. Can we treat in the same way the other finite Coxeter types?

Exercises III

III.1 Let r be the rank function of a matroid on E and suppose that e is a loop of this matroid. Prove that e is contained in every flat.

III.2 Let r be the rank function of a matroid on E and suppose that e is an isthmus of this matroid. Prove that the lattice of flats of $r_{/e}$ and $r_{\setminus e}$ are isomorphic.

III.3 Let $d, n \in \mathbb{N}$, $d \leqslant n$. The *uniform matroid* $U_{m,d}$ *of rank* d *on* n *elements* has ground set $[n]$ and rank function defined by

$$r(A) = \begin{cases} |A| & \text{if } |A| \leqslant d \\ d & \text{otherwise} \end{cases}$$

Prove that the dual to $U_{n,d}$ is $U_{n,n-d}$.

III.4 Let $\mathscr{A}$ be the arrangement of four hyperplanes H_1, H_2, H_3, H_4 in $\mathbb{R}^2$ with equations:

$$H_1 : x + y = 0, H_2 : x = 0, H_3 : y = 0, H_4 : x - y = 0$$

Draw a picture of this arrangement and of the Hasse diagram of $\mathcal{L}(\mathscr{A})$. Compute the Tutte polynomial of the associated matroid and verify the formula for the number of regions of $\mathscr{A}$ in this case.

III.5 Consider the graph

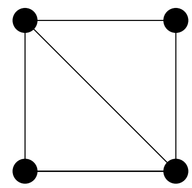

Draw the Hasse diagram of the lattice of flats of the associated matroid. Compute the Tutte polynomial and the number of regions of the associated graphic arrangement. Check it with a sketch of the arrangement (it is an arrangement in $\mathbb{R}^4$, but all hyperplanes contain the line ℓ in direction $(1,1,1)$, thus by quotienting by ℓ one can attempt a 3-dimensional sketch.

III.6 Let $\mathcal{L}$ be a finite geometric lattice. The *Möbius function* of $\mathcal{L}$ is the function $\mu : \mathcal{L} \to \mathbb{N}$ defined via the following recursion:

$$\mu(\widehat{0}) = 1, \quad \text{and} \sum_{y \leqslant x} \mu(y) = 0 \text{ for all } x \in \mathcal{L}_{>\widehat{0}}.$$

- Compute μ for the geometric lattice $\mathcal{L}$ in Example 3.1.2 and find an interpretation of the sum $\sum_{x \in \mathcal{L}} |\mu(x)|$ in this case.
- Compute the values of μ for the poset of Example 3.3.2 and consider the numbers you obtain by summing those values over every rank-level of the poset. Then compute the chromatic polynomial of the graph of the same example. What do you notice?
- Let now $\mathcal{L}$ be any geometric lattice. Prove that, for all $x \in \mathcal{L}$,

$$\mu(x) = \sum_{\substack{S \subseteq A(\mathcal{L}) \\ \bigvee S = x}} (-1)^{|S|}$$

where the sum is over all sets S of atoms of $\mathcal{L}$ whose join equals x. (Hint: prove that $\mu(x)$ and the expression on the r.-h.s. satisfy the same recursion).
- Let r be the rank function of a matroid with ground set E that has no loops. Show that

$$T_r(1-t, 0) = (-1)^{r(E)} \sum_{F \in \mathcal{L}_r} \mu(F) t^{r(E)-r(F)}.$$

[Hint: use the previous item in this exercise.]
- Prove that, for every arrangement of hyperplanes $\mathscr{A}$, the number of regions determined by $\mathscr{A}$ equals $\sum_{x \in \mathcal{L}(\mathscr{A})} |\mu(x)|$.

- Prove that for every graph G on n vertices the chromatic polynomial of G equals

$$\chi_G(t) = (-1)^{r(G)} t^{c(G)} \sum_{i=0}^{r(G)} \mu_i(\mathcal{L}_{r_G}) t^i$$

where for a geometric lattice $\mathcal{L}$ with rank function ρ we write $\mu_i(\mathcal{L}) = \sum_{\rho(x)=i} \mu(x)$ for the sum of the values of μ on all elements of $\mathcal{L}$ of rank i.

Check this formula with the lattice of Exercise III.5.

Chapter 4

Signed graphs

4.1 Definition

A loopless signed graph is, roughly speaking, a graph with two kinds of edges ("standard" and "half-" edges) and a "sign" associated to every edge.

Formally we will think of a tuple

$$\Sigma = (G, \sigma), \quad G = (V, E, h : E \to V, t : E_r \to V), \quad \sigma : E_r \to \{\pm\}$$

where $E_r \subseteq E$ is the set of all "regular" edges, while $E_{1/2} := E \setminus E_r$ is the set of half-edges. We will only need to consider loopless graphs, thus we assume $|V(e)| = 2$ for all $e \in E_r$.

Note that σ labels every edge with a sign $+$ or $-$. Therefore every trail has a sign as well, obtained as the product of the signs of all its edges. A circuit is called *balanced* if its sign $+$, unbalanced otherwise. A connected component of a signed graph is called *balanced* if every circuit in it is balanced, unbalanced otherwise. For every $A \subseteq E$ we call $b(A)$ the number of balanced connected components of the restriction of the signed graph on the set A (as in the unsigned case, this is the graph with same vertex set V but only the edges in A).

Remark 4.1.1. In particular, any cycle in a balanced component of a signed graph has positive sign. This is easy to see by noticing that a cycle decomposes as a disjoint union of circuits (each with positive sign) and "repeated" edges (and the product of two negatives is a positive).

Lemma 4.1.2. *Let* p_1, p_2 *be two paths with the same endpoints in a balanced component of a signed graph. Then the product of the signs of the edges of* p_1 *equals the product of the signs of the edges of* p_2.

Proof. The sign of the union of p_1 and p_2 is by definition the product of the signs of the two paths. But this union is a cycle which must have positive sign by the preceding remark. The claim follows. □

4.2 Coloring

Let $k \in \mathbb{N}$ and let $[[k]] := \{-k, -k+1, \ldots, 0, 1, \ldots, k\}$. A k-coloring of a loopless signed graph Σ is an assignment $\gamma : V \to [[k]]$ such that:

(i) $\gamma(h(e)) \neq \sigma(e)\gamma(t(e))$ for all $e \in E_r$,

(ii) $\gamma(h(e)) \neq 0$ for all $e \in E_{1/2}$.

As in the case of unsigned graphs, we call *precoloring* any function $V \to [[k]]$ and, given any $A \subseteq E$, we define the set of precoloring which fail to satisfy the constraints imposed by the edges in A:

$$\mathrm{Pre}_\Sigma(A) := \left\{ \gamma : V \to [[k]] \;\middle|\; \begin{array}{ll} \gamma(h(e)) = \sigma(e)\gamma(t(e)) & \text{for } e \in A \cap E_r, \\ \gamma(h(e)) = 0 & \text{for } e \in E_{1/2} \cap A \end{array} \right\}$$

Then any $\gamma \in \mathrm{Pre}_\Sigma(A)$ must have constant value across every connected component of $G[A]$. Moreover, notice that every unbalanced component of $G[A]$ must contain an unbalanced cycle, say with sequence of edges-vertices $v_0, e_1, v_1 \ldots, e_m, v_m = v_0$, and we have

$$\gamma(v_0) = \sigma(e_1)\gamma(v_1) = \ldots = \sigma(e_1)\cdots\sigma(e_m)\gamma(v_m) = -\gamma(v_0),$$

the last equality because $v_m = v_0$ and the cycle is unbalanced. In particular, $\gamma(v_0) = 0$ and thus γ must be 0 across every unbalanced component of A.

Since the (constant) value of γ on any balanced connected component of $G[A]$ can be chosen freely, we conclude that

$$|\mathrm{Pre}_\Sigma(A)| = (2k+1)^{b(A)}.$$

Just as in the case of unsigned graphs, applying inclusion-exclusion on the family of all subsets of E we obtain that the function $\chi_\Sigma(t)$ that for odd t counts the number of $\frac{(t-1)}{2}$-colorings of Σ can be written

$$\chi_\Sigma(t) = \sum_{A \subseteq E} (-1)^{|A|} t^{b(A)}$$

and is therefore a polynomial.

Coup de théâtre: we'll show next that this is in fact, up to a factor, the characteristic polynomial of a matroid!

4.3 Matroids

Definition 4.3.1. Let Σ be a (loopless) signed graph on n vertices. Define a function

$$r_\Sigma : 2^E \to \mathbb{N}, \qquad A \mapsto r_\Sigma(A) := n - b(A)$$

Theorem 4.3.2. *The function r_Σ is the rank function of a matroid on E.*

Proof. We will verify the rank axioms, but first a general observation.

G.O. When adding an edge to a signed graph Σ, the number of balanced components can drop in two ways: either ome balanced component becomes unbalanced (i.e., by adding a half-edge, by adding an edge that creates an unbalanced cycle or by joining a balanced component with an unbalanced one) or two balanced components are joined via the new edge to one bigger balanced component . In any case the number of balanced components either stays the same or it decreases at most by one, i.e,

$$r_\Sigma(A) \leqslant r_\Sigma(A \cup \{e\}) \leqslant r_\Sigma(A) + 1. \tag{4.1}$$

r0 Clearly r_Σ is never negative. In order to prove that $r_\Sigma(A) \leqslant |A|$ for all $A \subseteq E$, start with the graph Σ_0 on V and with no edges: here every vertex is its own (balanced!) component. Therefore $r_\Sigma(\emptyset) = n - n = 0$. Now add the edges in A to Σ_0 one by one. By the G.O., the number of balanced components decreases at most by one at each step, thus the rank increases at most by one at each step (second inequality in Equation (4.1)). Thus, $r_\Sigma(A) \leqslant |A|$ as required.

r1 Obviously it is enough to show that $r_\Sigma(A) \leqslant r_\Sigma(A \cup \{e\})$ for all $A \subseteq E$ and all e. But this is implied directly from the G.O. above (first inequality of Equation (4.1)).

r2 Here it is enough to show axiom r2 for the case where $X = A \cup \{e\}$ and $Y = A \cup \{f\}$, for any given $A \subseteq E$ and $e, f \in E$ (see Exercise IV.1). In this case we have to show

$$r_\Sigma(A \cup \{e, f\}) - r_\Sigma(A \cup \{f\}) \leqslant r_\Sigma(A \cup \{e\}) - r_\Sigma(A).$$

Now, if the right-hand side equals 1 then the inequality holds trivially, since by (r1) the left-hand side is at most 1.

Suppose then that $r_\Sigma(A \cup \{e\}) = r_\Sigma(A)$ – in particular, if e has one end in a balanced component, then it must be a regular edge with both ends in the same balanced component of A. By way of contradiction, suppose now that $r_\Sigma(A \cup \{e, f\}) - r_\Sigma(A \cup \{f\}) = 1$. The latter means that the number of balanced components changes by adding e to $A \cup \{f\}$. By the G.O., this means that either

(1) e connects two balanced components of $A \cup \{f\}$ – a contradiction because both ends of e are in the same component of A, and *a fortiori* of $A \cup f$, or

(2) e "closes" an unbalanced circuit C in a balanced component of $A \cup \{f\}$. But this can only happen if C also contains f. Now, if f

connects two balanced components of A, then e would, too – a contradiction. So the two ends of f lie in the same balanced component of A and (by the connectedness of this component) are connected by a path p whose sign agrees with that of f (by Lemma 4.1.2). Now we can replace f by p in C and obtain a closed, unbalanced cycle containing e in $A \cup e$, thus e "closes" an unbalanced cycle in a balanced component of $\Sigma[A]$, contradicting the assumption.

In all possible cases we reached a contradiction and the desired inequality is proved.

□

4.4 Arrangements from signed graphs

We now show that the matroid defined by r_Σ is in fact obtainable as the matroid of linear dependencies of a set of vectors.

Definition 4.4.1. Let Σ be a loopless signed graph on the vertex set V with set of edges E. For each edge $e \in E$ define a vector $m_e \in \mathbb{R}^V$ as follows:

$$(m_e)_v := \begin{cases} 1 & \text{if } v = h(e) \\ \sigma(e)(-1) & \text{if } v = t(e) \\ 0 & \text{otherwise.} \end{cases}$$

In particular, if $e \in E_{1/2}$ the vector m_e has only one nonzero component.
Let M_Σ the $V \times E$ matrix with columns m_e and, for every $A \subseteq E$, let $M_\Sigma[A]$ denote the matrix constructed with only the columns indexed by A.

The following theorem states that r_Σ is the rank function of the matroid obtained as in Theorem 2.2.15 from the vectors m_e where e ranges over E.

Theorem 4.4.2. *For every* $A \subseteq E$, $r_\Sigma(A) = \operatorname{rank} M_\Sigma[A]$.

The proof will use the following lemma.

Lemma 4.4.3. *Let* Σ *be a connected signed graph on* n *vertices. Then* $\operatorname{rank} M_\Sigma = n-1$ *if* Σ *is balanced, and* $\operatorname{rank} M_\Sigma = n$ *otherwise.*

Proof. Let T be a set of regular (non-half) edges that is a spanning tree of Σ. We can suppose without loss of generality that, if Σ has half-edges, then one half-edge is attached to a leaf v_0 of T. Then T has size $n-1$ and we can order the vertices of Σ beginning with v_0, and continuing in such a way that, for all i, the subgraph of T induced by the vertices $v_0, \dots, v_i$ is connected. Then the matrix M_Σ with the rows ordered this way can be written as follows (where the first row is only present in case v_0 has a half-edges attached to it)

$$\begin{pmatrix} (1) & \mp 1 & * & * & \cdots & * & * & \cdots & * \\ (0) & \pm 1 & * & & & * & & & * \\ (0) & 0 & \pm 1 & * & & * & & & * \\ (0) & 0 & 0 & \pm 1 & \ddots & * & & & * \\ \vdots & \vdots & & \ddots & \ddots & * & & & * \\ (0) & 0 & & \cdots & 0 & \pm 1 & * & \cdots & * \end{pmatrix}$$

and it has clearly rank at least $n-1$. Since M_Σ has n rows, its rank is at most n. Now, it is enough to show the following, which we leave as an exercise.

- If Σ is balanced, then the parenthesized column is not present and every m_e, for $e \notin T$, is a linear combination of the set $\{m_t \mid t \in T\}$. (Hint: for every such e consider the (unique) circuit contained in $T \cup \{e\}$.
- If Σ is not balanced, then either it has an half-edge, in which case the parenthesized column is present and so the matix has rank n, or there is no half-edge. In this last case, there is an $e \in E \setminus T$ that is linearly independent from $\{m_t \mid t \in T\}$.

□

Proof of Theorem 4.4.2. First notice that $M_\Sigma[A]$ is the matrix associated to the restriction of Σ to the edge-set A. Therefore, it is enough to prove the claim for $A = E$, i.e., we have to prove that $\operatorname{rank} M_\Sigma = r_\Sigma(E)$. Moreover, if Σ is disconnected we can let $V_1, \ldots, V_k$ and $A_1, \ldots, A_k$ be the vertex sets, resp. the edge sets of the connected components, and we see that the matrix M_Σ has block-diagonal form

$$\begin{array}{c} \\ V_1 \\ V_2 \\ \vdots \\ V_k \end{array} \begin{array}{c} \begin{array}{cccc} A_1 & A_2 & \cdots & A_k \end{array} \\ \begin{pmatrix} M_{\Sigma_1} & 0 & 0 & 0 \\ 0 & M_{\Sigma_2} & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & M_{\Sigma_k} \end{pmatrix} \end{array}$$

showing that $\operatorname{rank} M_\Sigma = \sum_i \operatorname{rank} M_{\Sigma_i}$ where Σ_i is the signed graph on the vertex set V_i with (signed) edges A_i. On the other hand, writing $n_i = |V_i|$ and setting $b_i = 1$ if Σ_i is balanced and $b_i = 0$ otherwise, we have that

$$\sum_i r_{\Sigma_i}(A_i) = \sum_i (n_i - b(A_i)) = \sum_i (n_i - b_i) = n - b(E) = r_\Sigma(E).$$

Therefore, it is enough to prove the claim for Σ connected. This has been done precisely in Lemma 4.4.3. □

The Tutte polynomial of the matroid r_Σ is of course related to the chromatic polynomial χ_Σ.

Lemma 4.4.4. *Let Σ be a loopless signed graph. Then*

$$\chi_\Sigma(t) = (-1)^n t^{b(E)} T_{r_\Sigma}(1-t, 0)$$

Proof. Direct computation:

$$T_{r_\Sigma}(1-t,0) = \sum_{A\subseteq E} (-t)^{(n-b(E))-(n-b(A))}(-1)^{|A|-(n-b(A))}$$

$$= (-1)^n t^{-b(E)} \sum_{A\subseteq E} (-1)^{|A|} t^{b(A)}$$

□

4.5 Arrangements associated to root systems

For every fixed dimension d there are only a finite number of finite (irreducible) groups of isometries of $\mathbb{R}^d$ generated by reflections. Those are the (irreducible) finite Coxeter groups, and their theory is a mainstay of algebra and geometry since its bases were laid by H.S.M. Coxeter[1] in the mid- 20. Century.

The classification of (irreducible) finite Coxeter groups identifies in every dimension one representative of each of four infinite families (indexed by A_n, B_n, C_n, D_n) and, in some (low) dimensions, a finite number of sporadic cases. One way to specify one of these groups, say G, is to give all isometries $f \in G$ that are reflections[2] - by definition, these generate the group G. Now, a (orthogonal) reflection f is fully specified by its fixed space $\mathrm{fix}(f) = \{x \in \mathbb{R}^d \mid f(x) = x\}$. Note that if f is a reflection, then $\mathrm{fix}(f)$ is a hyperplane in $\mathbb{R}^d$. Therefore, G is given by the associated "arrangement of reflection hyperplanes". The arrangements that arise this way are explicitly listed, e.g., by means of a choice of their normal vectors. The set of normal vectors is usually referred to as the set of *(positive) roots* of the corresponding Coxeter group[3]. (The word "positive" is there because root systems contain the opposite of each root, but opposite vectors define the same hyperplane.) Notice that the group acts freely and transitively on the set of regions of the corresponding arrangement, so there are as many regions as there are group elements.

We list here a set of normal vectors for the hyperplane arrangements of the four infinite families. We use the notation ϵ_i to denote the i-th standard basis

[1] https://en.wikipedia.org/wiki/Harold_Scott_MacDonald_Coxeter
[2] A reflection
[3] See for instance Bourbaki

vector in $\mathbb{R}^n$, so ϵ_i has all components equal to 0 except the i-th component being equal to 1.

A_{n-1}[4]: The vectors $\epsilon_i - \epsilon_j$ in $\mathbb{R}^n$ for $n \geqslant i > j \geqslant 1$

B_n: The vectors $\epsilon_i - \epsilon_j$ as well as $\epsilon_i + \epsilon_j$ and ϵ_i in $\mathbb{R}^n$ for $n \geqslant i > j \geqslant 1$

C_n: The vectors $\epsilon_i - \epsilon_j$ as well as $\epsilon_i + \epsilon_j$ and $2\epsilon_i$ in $\mathbb{R}^n$ for $n \geqslant i > j \geqslant 1$

D_n: The vectors $\epsilon_i - \epsilon_j$ and $\epsilon_i + \epsilon_j$ in $\mathbb{R}^n$ for $n \geqslant i > j \geqslant 1$

We first notice that the arrangement of type B and C are equal (since ϵ_i and $2\epsilon_i$ define the same hyperplane). Then, we notice that the set of vectors in types A, B, D can be obtained as vectors of the type m_e for edges of suitable signed graphs, that we will call $\Sigma(A_{n-1})$, $\Sigma(B_n)$, $\Sigma(D_n)$ and we describe below.

$\Sigma(A_{n-1})$ is the graph K_n^+ obtained from the complete graph on n vertices by assigning a positive sign to each vertex.

$\Sigma(B_n)$ is the graph on n vertices with a pair of opposite-signed edges between every pair of vertices, and a half-edge at every vertex.

$\Sigma(D_n)$ is the graph on n vertices with a pair of opposite-signed edges between every pair of vertices.

Now the number of regions in each such arrangement is, by the general principle, given by the Tutte polynomial of the associated matroid r_Σ – and, by Lemma 4.4.4 as

$$T_{r_\Sigma}(2, 0) = (-1)^{n-b(E)}\chi_\Sigma(-1).$$

In order to compute this number it is now enough to compute the polynomial $\chi_\Sigma(t)$ for the signed graphs associated to the different root systems.

A_{n-1} The polynomial is $\chi_{\Sigma(A_{n-1})}(t) = \prod_{i=0}^{n-1}(t - i)$
The constraints are as in the (unsigned) complete graph: given t colors we can assign any of them to the first vertex, only $t - 1$ to the second, and so on.

Thus the number of regions of the arrangement is $n!$ (this agrees with the fact that the reflection group of type A_{n-1} is the symmetric group S_n of all permutations of n elements - and of those there are notoriously $n!$.

B_n The polynomial is $\chi_{\Sigma(B_n)}(t) = \prod_{i=1}^{n}(t - 2i + 1)$
Here the constraints are as follows: between any two vertices we have two signed edges, that forbid the two vertices to have the same or the opposite color. Moreover the half-edges forbid the color 0

everywhere. So that for (odd) t signed colors we have $(t-1)$ possibilities for vertex 1, $(t-3)$ for vertex 2 and so on.

Thus the number of regions of the arrangement (and hence the order of the reflection group) is $2n(2n-2)(2n-4)\cdots = 2^n n!$

D_n The polynomial is $\chi_{\Sigma(D_n)}(t) = (t-n+1)\prod_{i=1}^{n-1}(t-2i+1)$
Here the constraints are as in B_n but 0 may be allowed: between any two vertices we have two signed edges, that forbid the two vertices to have the same or the opposite color. Therefore we can sum over the possibilities: either 0 is not used, or it is used once, and there are n possibilities to assign it, while all other vertices must be colored differently, and according to the rules of B_n:

$$\chi_{\Sigma(D_n)}(t) = \prod_{i=1}^{n}(t-2i+1) + n\prod_{i=1}^{n-1}(t-2i+1).$$

The number of regions of the arrangement (and hence the order of the reflection group) is $n(2n-2)(2n-4)\cdots = 2^{n-1}n!$

Exercises IV

IV.1 Prove that axiom (r2) in the definition of a matroid rank function is equivalent to the statement that

$$r(A \cup \{e\}) + r(A \cup \{f\}) \geqslant r(A) + r(A \cup \{e, f\}) \quad \text{for all } A \subseteq E,\ e, f \in E.$$

IV.2 Let Σ be a connected, loopless signed graph and let T be any spanning tree of Σ (i.e., T is an acyclic set of regular edges that are incident to every vertex). Prove:

- For every regular $e \in E_r \setminus T$ there is a unique circuit contained in $T \cup \{e\}$. We call this circuit $C(T, e)$.
- If Σ is unbalanced and has no half-edges, then there is an edge $e \in E$ such that $C(T, e)$ is unbalanced.

Chapter 5

A polyhedral interlude

5.1 Motivation: a loose end about flows on graphs

Recall that an integral k-flow on a graph $G = (V, E, h, t)$ is an assignment $f : E \to \{-k, \dots, k\}$ such that

$$\sum_{h(e)=v} f(e) = \sum_{t(e)=v} f(e) \quad \text{(as an equation in } \mathbb{Z}\text{) for all } v \in V. \tag{5.1}$$

The task we set ourselves in Section 1.3 was to determine the function $\varphi_G^{\mathbb{N}}$ that counts the number of nowhere-zero integral k-flows on G.

Definition 5.1.1. Call an integer k- flow on G *positive* if $f(E) > 0$. Let $I_G(k)$ be the number of positive integral k-flows on G.

We can make any flow positive by "flipping" edges, i.e., performing the following operation: Given a graph $G = (V, E, h, t)$ and $A \subseteq E$, let

$$G^{\neg A} := (V, E, h', t'), \quad \text{where } h'(a) = t(a) \text{ and } t'(a) = h(a) \text{ for all } a \in A.$$

Intuitively, as the word says, we "flip" head and tail of all edges in A.

Lemma 5.1.2. *If f is a nowhere-zero flow on G, let $A := \{e \in E \mid f(e) < 0\}$. Then the assignment*

$$f'(e) := \begin{cases} -f(e) & \textit{if } e \in A \\ f(e) & \textit{otherwise} \end{cases}$$

is a positive integer k-flow on $G^{\neg A}$.

Proof. We only have to show that f' is an integer k-flow. □

In particular, for every nowhere-zero integral flow there is exactly one set $A \subseteq E$ such that f' is a positive integral flow on $G^{\neg A}$. This shows that

$$\varphi_G^{\mathbb{N}}(k) = \sum_{A \subseteq E} I_{G^{\neg A}}(k)$$

and therefore we are led to consider the function $I_G(k)$.

Recall the incidence matrix of the graph G, a matrix $A \in \mathbb{R}^{V \times E}$ whose v-th component in the e-th column is

$$(A_e)_v = \begin{cases} 1 & \text{if } v = h(e) \\ -1 & \text{if } v = t(e) \\ 0 & \text{otherwise.} \end{cases}$$

Now a positive k-flow on G is any $f \in \mathbb{Z}^E$ such that

$$Af = 0, \qquad \mathbf{0} < f < \mathbf{k}. \tag{5.2}$$

where $\mathbf{0} = (0, \ldots, 0)$, $\mathbf{k} = (k, \ldots, k)$ and the inequalities are understood componentwise.

Now consider the set

$$U := \{x \in \mathbb{R}^E \mid Ax = 0, \mathbf{0} < x < \mathbf{1}\}$$

and its closure $\overline{U} := \{x \in \mathbb{R}^E \mid Ax = 0, \mathbf{0} \leqslant x \leqslant \mathbf{1}\}$.

Then we have

$$I_G(k) = \#\mathbb{Z}^E \cap kU.$$

The fact that $\overline{U}$ is a convex polytope with integer vertices opens up the possibility to compute $I_G(k)$ using Ehrhart theory.

5.2 Polyhedra and polytopes

Recall that, when writing inequalities between vectors we intend the inequalities to be taken componentwise. Moreover, in this section A will denote any $n \times m$ matrix with real entries.

Definition 5.2.1. A **polyhedron** is any subset of $\mathbb{R}^n$ of the form

$$P = \{x \in \mathbb{R}^n \mid Ax \leqslant b\}$$

for some matrix $A \in \mathbb{R}^{m \times n}$ and some vector $b \in \mathbb{R}^m$.

Example 5.2.2. In Figure 1 we depict two instances of polyhedra in $\mathbb{R}^2$.

Definition 5.2.3. A **polytope** is any subset of $\mathbb{R}^n$ of the form

$$P = \operatorname{conv}\{x_1, \ldots, x_k\} := \{\sum \lambda_i x_i \mid \lambda_i \in \mathbb{R}_{\geqslant 0}, \sum \lambda_i = 1\}$$

for some finite subset $\{x_1, \ldots, x_k\} \subseteq \mathbb{R}^n$.

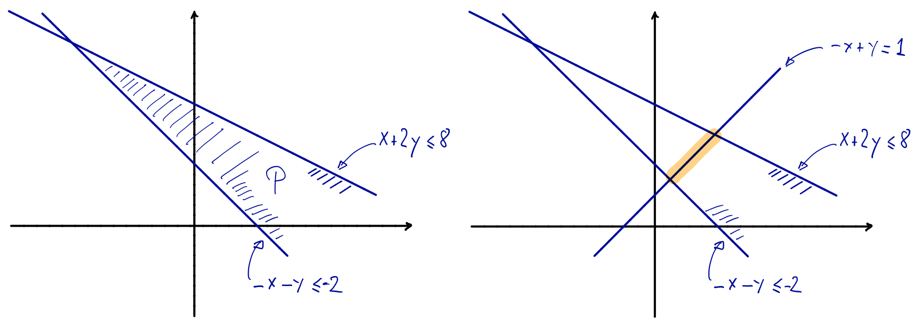

(a) The (unbounded) polyhedron

$$\left\{ \begin{bmatrix} 1 & 2 \\ -1 & -1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \leqslant \begin{bmatrix} 8 \\ -2 \end{bmatrix} \right\}$$

(b) The (bounded) polyhedron

$$\left\{ \begin{bmatrix} 1 & 2 \\ -1 & -1 \\ -1 & 1 \\ 1 & -1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} \leqslant \begin{bmatrix} 8 \\ -2 \\ 1 \\ -1 \end{bmatrix} \right\}$$

Figure 1

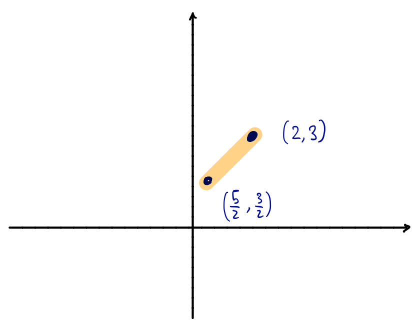

Figure 2: The polytope conv $\left\{(2,3), (\frac{5}{2}, \frac{3}{2})\right\}$

Example 5.2.4. In Figure 2 we depict a polytope in $\mathbb{R}^2$, identical to the bounded polyhedron of Figure 1b.

The goal of this section will be to prove the following fundamental theorem of polytope theory.

Theorem 5.2.5. *A subset of $\mathbb{R}^n$ is a polytope if and only if it is a bounded polyhedron.*

The next statement is our stepping stone towards proving Theorem 5.2.5.

Theorem 5.2.6. *Let $a_1, \dots, a_m, v \in \mathbb{R}^n$. Then either*

(1) $v = \lambda_1 a_1 + \dots + \lambda_m a_m$ *with* $\lambda_i \geqslant 0$*, or*

(2) there is $c \in \mathbb{R}^n$ *such that the hyperplane* $c^\perp$ *contains* $t-1$ *linearly independent vectors from the set* $\{a_1, \dots, a_m\}$ *such that* $c \cdot v < 0$ *and* $c \cdot a_1, \dots, c \cdot a_m \geqslant 0$*, where* $t = \dim\langle a_1, \dots, a_m, v\rangle$.

Proof. First, notice that (1) and (2) are mutually exclusive: if both hold, then we would have $0 > c \cdot v = \lambda_1 c \cdot a_1 + \dots + \lambda_m c \cdot a_m \geqslant 0$, a contradiction. Thus, it will suffice to show that in any case at least one of the options holds.

Let $V := \langle a_1, \dots, a_m\rangle$. If $v \notin V$, then $V \subsetneq \mathbb{R}^n$ is a proper subspace, and thus it is contained in a hyperplane H. Choose c to be any normal vector to H. We have $c \cdot v \neq 0$ and so, after possibly switching sign of c, we have $c \cdot v > 0$. Moreover, as $V \subseteq H$, we also have $c \cdot a_1 = \dots = c \cdot a_m = 0$, so case (2) holds.

Assume from now $v \in V$ and, without loss of generality, $V = \mathbb{R}^n$, so that $t = n$. Choose then a linearly independent set $\{a_{i_1}, \dots, a_{i_n}\}$ among the a_i and set

$$D_0 := \{i_1, \dots, i_n\}$$

and for all $i \geqslant 0$ iterate the following procedure.

(i) Consider the (unique) expression $v = \sum_{j \in D_i} \lambda_j a_j$. If $\lambda_j \geqslant 0$ for all $j \in D$, we are in case (1) and we stop.

(ii) Otherwise, choose $h := \min\{j \in D \mid \lambda_j < 0\}$ and let $H = \langle a_j \mid j \in D \setminus \{h\}\rangle$, a hyperplane of which we can choose a normal c with $c \cdot a_h = 1$, so that $c \cdot v = \lambda_h < 0$.

(iii) If $C \cdot a_i \geqslant 0$ for all $i = 1, \dots, m$, we are in case (2) and we can stop.

(iv) Otherwise, let $s := \min\{j \in [m] \mid c \cdot a_j < 0\}$, set

$$D_{i+1} := (D_i \setminus \{h\}) \cup \{s\}$$

and repeat.

We must prove that this iteration terminates; by way of contradiction, assume that this is not the case. Then, since there is only a finite number of options for choosing $D_i \subseteq [m]$, we must have $D_k = D_l$ for some $k < l$.
Now at every step i between k and l one element is removed from (and one is added to) D_i. Let r be the maximum among all elements that are removed at some step i, $k \leqslant i < l$ and let p be a step where this happens (so $D_{p+1} = D_p \setminus \{r\} \cup \{j\}$ for some j). Moreover, since $D_k = D_l$ there must be a $q \neq p$, $k \leqslant q < l$, such that r is added to D_q at the q-th step. In particular,

$$D_p \cap \{r+1, \dots, m\} = D_q \cap \{r+1, \dots, m\} \tag{5.3}$$

Letting c_q be the vector from item (ii) in the q-th iteration and considering the (unique) expression $v = \sum_{j \in D_p} \lambda_j a_j$ we reach the desired contradiction:

$$0 > c_q \cdot v = c_q \cdot \Big(\sum_{j \in D_p} \lambda_j a_j \Big) = \sum_{j \in D_p} \lambda_j c_q \cdot a_j > 0.$$

the first inequality is by the choice of c_q in item (ii). The second inequality holds because, in fact, every summand is nonnegative and the r-th is positive. To prove this, remember that

$*$ since r is removed from D_p, by (ii) $r = \min\{j \in D_p \mid \lambda_j < 0\}$, and

$\star$ since r is added to D_q by (iv) $r = \min\{j \in [m] \mid c_q \cdot a_j < 0\}$.

Now let $j \in D_p$:

If $j < r$ then $\lambda_j \geqslant 0$ by $*$ and $c_q \cdot a_j \geqslant 0$ by $\star$, thus $\lambda_j c_q \cdot a_j \geqslant 0$.

If $j = r$ then $\lambda_j < 0$ by $*$ and $c_q \cdot a_j < 0$ by $\star$, thus $\lambda_j c_q \cdot a_j > 0$.

If $j > r$ then $j \in D_q$ by Equation (5.3) and since the element removed from D_q is smaller than r (by maximality in the choice of r), then the definition of c_q implies $c_q \cdot a_j = 0$.

□

Let us review the two alternatives in the claim of Theorem 5.2.6. Claim (1) can be rephrased by saying that "v is in the convex cone generated by $a_1, \ldots, a_m$". In fact, a *convex cone* is any $C \subseteq \mathbb{R}^n$ such that $\lambda x + \mu y \in C$ for all $x, y \in C$ and all $\lambda, \mu \geqslant 0$, and the smallest convex cone containing a given set $x_1, \ldots x_m \in \mathbb{R}^n$ is called the "cone generated by $x_1, \ldots x_m$ and written

$$\text{cone}\{x_1, \ldots, x_m\} := \{\lambda_1 x_1 + \ldots + \lambda_m x_m \mid \lambda_i \geqslant 0\}$$

and any convex cone arising this way is called *finitely generated*.

On the other hand we call *polyhedral cone* any nonempty $C \subseteq \mathbb{R}^n$ of the form

$$C = \{x \in \mathbb{R}^n \mid Ax \leqslant 0\} \tag{5.4}$$

for some matrix A.

Example 5.2.7. The cone in Figure 3c is generated, e.g., by the vectors $\begin{bmatrix} 2 \\ -1 \end{bmatrix}$ and $\begin{bmatrix} 1 \\ -1 \end{bmatrix}$.

The alternative expressed in Theorem 5.2.6 then can be understood to express the fact that any point v that is not in the convex cone generated by $a_1, \ldots, a_m$ is also not in the smallest polyhedral cone containing the a_is. In this vein, and more precisely, we have the following.

Proposition 5.2.8. *A convex cone is polyhedral if and only if it is finitely generated.*

Proof.
Finitely generated ⇒ polyhedral. Let $X := \{x_1, \dots, x_m\} \subset \mathbb{R}^n$ and assume without loss of generality that $x_1, \dots, x_n$ span $\mathbb{R}^n$. Every nonzero vector $c \in \mathbb{R}^n$ defines a halfspace $\{x \mid c \cdot x \leqslant 0\}$ bounded by the hyperplane $\{x \mid c \cdot x = 0\}$, and we consider the set $\mathcal{H}$ of all halfspaces containing X and whose bounding hyperplanes are spanned by $n-1$ linearly independent elements among the x_i. For each halfspace $H \in \mathcal{H}$ choose a defining vector c_H. By Theorem 5.2.6, we have

$$\operatorname{cone}\{x_1, \dots, x_m\} = \bigcap \mathcal{H}$$

and since there are finitely many elements in $\mathcal{H}$, the cone is polyhedral (with respect to the matrix whose rows are the vectors c_H, $H \in \mathcal{H}$).

Polyhedral ⇒ finitely generated. Consider a polyhedral cone $C := \{Ax \leqslant 0\}$ and call $a_1, \dots, a_m$ the rows of the matrix A. Then the cone generated by the $a_1, \dots, a_m$ is polyhedral by the previous item ("the other direction"), and so there are $b_1, \dots, b_t \in \mathbb{R}^n$ such that

$$\operatorname{cone}\{a_1, \dots, a_m\} = \{Bx \leqslant 0\}$$

for the matrix B whose rows are the b_is.

Claim. $\operatorname{cone}\{b_1, \dots, b_t\} = C$, so C is finitely generated.

Proof. We call $K := \operatorname{cone}\{b_1, \dots, b_t\}$ and split the proof in two parts. First we prove $K \subseteq C$. In fact, by definition $a_i \in K$ and so $b_j \cdot a_i \leqslant 0$ for all i, j. Thus $Ab_j \leqslant 0$, and consequently $b_j \in C$, for all j.
Now we prove $C \subseteq K$. If there was a $y \in C \setminus K$, then by Theorem 5.2.6 there would be c with $c \cdot y < 0$ and $c \cdot b_i \geqslant 0$ for all i. But then we would have

$$(-c) \cdot b_i \leqslant 0 \text{ for all } i, \text{ so } -c \in \operatorname{cone}\{a_1, \dots, a_m\},$$

and we could write $-c = \sum_i \lambda_i a_i$ with all $\lambda_i \geqslant 0$. In particular, for every $x \in C$ we'd have

$$(-c) \cdot x = (\sum_i \lambda_i a_i) x = \sum \lambda_i a_i \cdot x \leqslant 0$$

the last inequality because $\lambda_i \geqslant 0$ and, since $x \in C$, $a_i \cdot x \leqslant 0$. In particular, since $y \in C$ we would have $c \cdot y \geqslant 0$, contradicting $c \cdot y < 0$.

□

The last preparation for the proof of Theorem 5.2.5 is a statement characterizing (possibly unbounded) polyhedra. Recall that we add subsets of $\mathbb{R}^n$ pointwise – i.e., for $A, B \in \mathbb{R}^n$ we write $A + B = \{a + b \mid a \in A,\ b \in B\}$.

Proposition 5.2.9. *Let $P \subseteq \mathbb{R}^n$. The following are equivalent.*

(1) P *is a polyhedron,*

(2) $P = Q + C$ *for a polytope* Q *and a polyhedral cone* C.

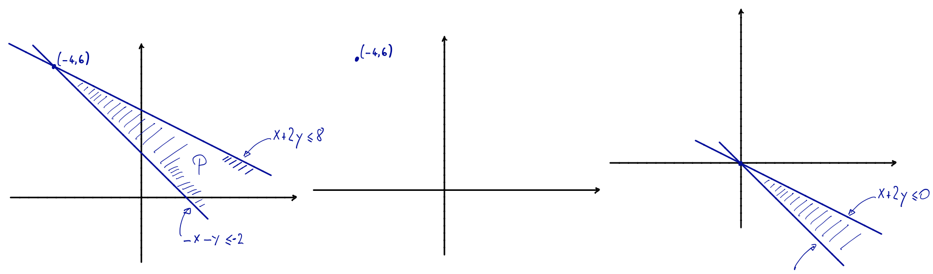

(a) (b) (c)

Figure 3: The polyhedron P in (a) can be written as $P = Q + C$ for the polytope Q from (b) and the polyhedral cone C from (c).

Proof of Proposition 5.2.9.

(2)⇒(1) Suppose $P = Q + C$ with Q a polytope, say $Q = \mathrm{conv}\{x_1, \dots, x_m\}$, and C a polyhedral cone, say $C = \mathrm{cone}\{y_1, \dots, y_t\}$. Then, $v \in P$ if and only if

$$\begin{pmatrix} v \\ 1 \end{pmatrix} \in \mathrm{cone}\left\{ \begin{pmatrix} x_1 \\ 1 \end{pmatrix}, \dots \begin{pmatrix} x_m \\ 1 \end{pmatrix}, \begin{pmatrix} y_1 \\ 0 \end{pmatrix}, \dots \begin{pmatrix} y_t \\ 0 \end{pmatrix} \right\} \tag{5.5}$$

Call K the finitely generated cone in Equation (5.5). By Proposition 5.2.8, K is polyhedral, i.e., there is a matrix A' such that

$$K = \left\{ A' \begin{pmatrix} v \\ \lambda \end{pmatrix} \leqslant 0 \right\} = \{Ax + \lambda \cdot b \leqslant 0\}$$

where b is the last column of A' and A is A' without the column b.

Therefore, $v \in P$ if and only if $Av \leqslant -b$, and thus $P = \{Ax \leqslant -b\}$ is a polyhedron.

(1)⇒(2) Suppose now that $P = \{Ax \leqslant b\}$ for some matrix A and vector b. Consider the polyhedral cone

$$K := \left\{ \left[\begin{array}{c|c} 0 & -1 \\ \hline A & -b \end{array} \right] \begin{pmatrix} x \\ \mu \end{pmatrix} \leqslant 0 \right\} = \left\{ \begin{pmatrix} x \\ \mu \end{pmatrix} \middle| \, Ax - \mu b \leqslant 0, \mu \geqslant 0 \right\}$$

By Proposition 5.2.8 this cone is finitely generated, say by vectors $\begin{pmatrix} x_1 \\ \mu_1 \end{pmatrix}, \dots, \begin{pmatrix} x_m \\ \mu_m \end{pmatrix}$, where we can suppose after rearranging and rescaling that there is $1 \leqslant k \leqslant m$ such that $\mu_i = 1$ for all $i \leqslant k$ and $\mu_i = 0$ for $i > k$.

Now by definition $P = \left\{ \begin{pmatrix} x \\ \mu \end{pmatrix} \in K \middle| \mu = 1 \right\}$, so $x \in P$ if and only if $\begin{pmatrix} x \\ 1 \end{pmatrix} \in K$, i.e., $\begin{pmatrix} x \\ 1 \end{pmatrix}$ can be expanded as a positive combination of the generators of K:

$$\begin{cases} x = & \sum_{i \leqslant k} \lambda_i x_i \\ 1 = & \lambda_1 + \ldots + \lambda_k \end{cases}, \quad \lambda_i \geqslant 0 \text{ for all } i,$$

Recalling the definition of cone and convex hull, this is equivalent to

$$x \in \operatorname{conv}\{x_1, \dots, x_k\} + \operatorname{cone}\{x_{k+1}, \dots, x_m\},$$

and the claim follows with $Q = \operatorname{conv}\{x_1, \dots, x_k\}$ and $C = \operatorname{cone}\{x_{k+1}, \dots, x_m\}$, the latter being polyhedral by Proposition 5.2.8.

□

Proof. Proof of Theorem 5.2.5 The claim follows from Proposition 5.2.9 by noting that the only bounded polyhedral cone is the cone generated by the single vector 0, i.e., a single point. □

5.3 Faces of polyhedra

Definition 5.3.1. A *face* of a polyhedron $P = \{Ax \leqslant b\}$ is any subset of P of the form

$$F = \{x \in P \mid A'x = b'\}$$

where $A'x \leqslant b'$ is a subsystem of (i.e., consists only of some of the inequalities of) $Ax \leqslant b$.

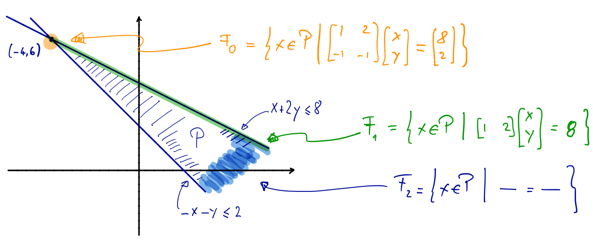

Three faces of the polytope P: a vertex (F_0), a facet (F_1) and the (improper) face $F_2 = P$.

5.3.1 Supporting hyperplanes

An equivalent characterization of faces of polyhedra uses the notion of "supporting hyperplane". We call a hyperplane H a supporting hyperplane for a polyhedron P if "P touches H but lies wholly on one side of H".

Definition 5.3.2. Let $P = \{Ax \leqslant b\}$ be a polyhedron. An affine hyperplane $H = \{c{\cdot}x = \delta\}$ is called a *supporitng hyperplane* for P if $\delta = \max\{c{\cdot}x \mid x \in P\}$.

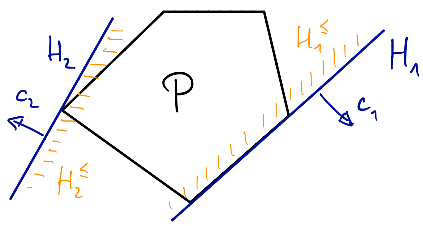

H_1 and H_2 are supporting hyperplanes of the polyhedron P.

Remark 5.3.3. Clearly every (nonempty) face F of a polyhedron P is of the form $F = P \cap H$ for some supporting hyperplane H. Indeed, suppose $F = \{x \in P \mid A'x = b'\}$ is nonempty. Then, if c is the sum of all rows of A', clearly F is the set of all points $x \in P$ attaining the maximum $\delta = \max\{c{\cdot}x \mid x \in P\}$.

Proving the converse to Remark 5.3.3 – that every intersection of P with a supporting hyperplane is a face – needs more work, based on a corollary of Theorem 5.2.6 that bears considerable importance in its own right.

Lemma 5.3.4 ("Farkas' Lemma")**.** *Let A be a matrix and b a vector. Then there exists $x \geqslant 0$ such that $Ax = b$ if and only if $yb \geqslant 0$ for each row vector y with $yA \geqslant 0$.*

Proof. One direction is clear: if $x \geqslant 0$, $Ax = b$ and $yA \geqslant 0$ then clearly $yb = yAx \geqslant 0$.

For the other direction, let $a_1, \dots, a_m$ be the columns of A and argue by contraposition: if there is no $x \geqslant 0$ with $Ax = b$, then $b \notin \operatorname{cone}\{a_1, \dots, a_m\}$ and thus, by Theorem 5.2.6, there is y with $yb < 0$ and $yA \geqslant 0$ (the latter being equivalent to $ya_i \geqslant 0$ for all i). $\square$

Corollary 5.3.5. *Let A be a matrix and b be a vector. Then the system $Ax \leqslant b$ has a solution x if and only if $yb \geqslant 0$ for each row vector $y \geqslant 0$ with $yA = 0$.*

Proof. The system $Ax \leqslant b$ has a solution if, and only if, the system

$$[I \mid A \mid -A] \begin{bmatrix} z_0 \\ \hline z_1 \\ \hline z_2 \end{bmatrix} = b$$

has a nonnegative solution. (Indeed: given a solution x to the first system, define vectors x^+ and x^- as the positive and negative part[1] of x, so that $x^+ \geqslant 0$, $x^- \leqslant 0$ and $x = x^+ + x^-$. then the (nonnegative) vector $\left[\begin{array}{c} b - Ax \\ \hline x^+ \\ \hline -x^- \end{array}\right]$ solves the second. Conversely, if a nonnegative $\left[\begin{array}{c} z_0 \\ \hline z_1 \\ \hline z_2 \end{array}\right]$ solves the second system, then $A(z_1 - z_2) = b - z_0 \leqslant 0$.) Now apply Farkas' Lemma to conclude that the latter system has a nonnegative solution if and only if $yb \geqslant 0$ for all y with $y[I \mid A \mid -A] = [yI \mid yA \mid -yA] \geqslant 0$ - i.e, $y \geqslant 0$ (from the first block) and $yA = 0$ (from the second and third block). ☐

We are ready now to state and prove the following result, known as "duality theorem for linear programming".

Proposition 5.3.6. *Let* A *be a matrix and* b, c *be vectors. Then*

$$\max\{c \cdot x \mid Ax \leqslant b\} = \min\{yb \mid y \geqslant 0,\ yA = c\}$$

whenever both sets are nonempty.

Proof. If $Ax \leqslant b$, $y \geqslant 0$ and $yA = c$, then $c \cdot x = yAx \leqslant yb$. This implies the inequality $\max \leqslant \min$, if the maxima and minima exist.

For the other inequality we need to show that there exist x, y with $Ax \leqslant b$, $y \geqslant 0$, $yA = c$ and $c \cdot x \geqslant yb$. This requirement we can translate in proving existence of a solution to

$$\underbrace{\begin{bmatrix} 0 & -1 \\ A & 0 \\ -c & b^t \\ 0 & A^t \\ 0 & -A^t \end{bmatrix}}_{=:\widetilde{A}} \begin{bmatrix} x \\ y^t \end{bmatrix} \leqslant \underbrace{\begin{bmatrix} 0 \\ b \\ 0 \\ c^t \\ -c^t \end{bmatrix}}_{\widetilde{b}}. \tag{5.6}$$

Now Corollary 5.3.5 implies that the solvability of (5.6) is equivalent to proving $w\widetilde{b} \geqslant 0$ for every $w \geqslant 0$ such that $w\widetilde{A} = 0$. In order to check this last statement write $w = (w_0, w_1, w_2, w_3, w_4)$ according to the block-decomposition in (5.6). Then the statement's premise is that $w_i \geqslant 0$ for all i and

$$(1) \quad w_1 A - w_2 c = 0, \qquad (2) \quad w_2 b^t + (w_3 - w_4) A^t \geqslant 0, \tag{5.7}$$

and we need to prove that

$$(3) \quad w_1 b + (w_3 - w_4) c^t \geqslant 0.$$

Now we distinguish two cases:

[1]Explicitly: $x_i^+ = \begin{cases} x_i & \text{if } x_i > 0 \\ 0 & \text{else.} \end{cases}$, $x_i^- = \begin{cases} x_i & \text{if } x_i < 0 \\ 0 & \text{else.} \end{cases}$

$w_2 > 0$. Here we write

$$w_2 w_1 b = w_2 b^t w_1^t \overset{(2)}{\geqslant} (w_3 - w_4) A^t w_1^t \overset{(1)}{=} (w_3 - w_4) w_2 c = w_2 (w_3 - w_4) c,$$

and (3) follows.

$w_2 = 0$ Since by assumption the sets over which maximum and minimum are taken in the proposition's claim are nonempty, we can choose elements x_0, y_0 with $Ax_0 \leqslant b$ and $y_0 \geqslant 0$, $y_0 A = c$. Then

$$w_1 b \geqslant w_1 A x_0 \overset{(1)}{=} 0 \overset{(2)}{\geqslant} (w_4 - w_3) A^t y_0^t = (w_4 - w_3) c^t$$

and (3) follows, concluding the proof.

□

Corollary 5.3.7. *Let* P *be a polyhedron. Then* $F \subsetneq P$ *is a face of* P *if and only if* $F = P \cap H$, *where* H *is a supporting hyperplane of* P.

Proof. Let $P = \{Ax \leqslant b\}$ and let $F = P \cap H$ with $H = \{c \cdot x = \delta\}$ and $\delta := \min\{c \cdot x \mid x \in P\}$, for some $c \neq 0$. Now by Proposition 5.3.6 we have $\delta = \min\{yb \mid y \geqslant 0, yA = c\}$, and we can choose y_0 that attains this minimum. Now consider $A'x \leqslant b'$, the system obtained from the rows of $Ax \leqslant b$ corresponding to positive entries of y_0. Then we have $F = \{x \in P \mid A'x = b'\}$ since for every $x \in P$ we have $Ax \leqslant b$ and so $c \cdot x = \delta$ is equivalent to $y_0 A x = y_0 b$ and in turn (erasing trivial rows) to $A'x = b'$.

The other direction is proved in Remark 5.3.3. □

5.3.2 Facets

Definition 5.3.8. A *facet* of a polyhedron P is any maximal proper face of P.

Given a polyhedron $P = \{Ax \leqslant b\}$ fix an enumeration $a_1, \dots, a_m$ of the rows of A (resp. the components of b), so that $P = \bigcap_{i=1,\dots m} \{a_i \cdot x \leqslant b_i\}$. Such an inequality $a_j \cdot x \leqslant b_j$ is called an *implicit equality* if $P \subseteq \{a_j \cdot x = b_j\}$. We can then let $A^{=} x \leqslant b^{=}$ denote the system of all implicit equalities and $A^{+} x \leqslant b^{+}$ the system consisting of all other inequalities, called *effective* inequalities.

Remark 5.3.9. Notice that any face $F = \{x \in P \mid A'x = b'\}$ can be written as $F = \{x \in P \mid {A^{+}}'x = {b^{+}}'\}$, where ${A^{+}}'x \leqslant {b^{+}}'$ consists of the effective inequalities in $A'x \leqslant b'$. (in fact, any implicit equality is... implicit in the requirement "$x \in P$).

Remark 5.3.10. If a system has an effective inequality $a_i \cdot x \leqslant b_i$ then there must be a point $x \in P \setminus \{a_i \cdot x = b_i\}$ and in particular this point x satisfies $A^{=} x = b^{=}$ and $A^{+} x < b^{+}$.

Recall that an inequality in a system of constraints is called *redundant* if its removal does not change the set of solutions.

Lemma 5.3.11. *Suppose* $a_i \cdot x \leqslant b_i$ *is an effective, non redundant inequality. Then* $F := \{x \in P \mid a_i \cdot x = b_i\}$ *is a facet of* P

Proof. Call $A'x \leqslant b'$ the system given by all effective inequalities other than $a_i \cdot x \leqslant b_i$. It is enough to prove that there is a point x_0 with

$$A^{=}x_0 = b^{=},\ A'x_0 < b'\ a_i \cdot x_0 = b_i$$

such an x_0 does not satisfy with equality any other row of $A'x \leqslant b'$, thus the only face bigger than F is P itself.

In order to find x_0 first recall Remark 5.3.10, which guarantees that we can find some x_1 with

$$A^{=}x_1 = b^{=} \text{ and } A^{+}x_1 < b^{+}$$

and, since $a_i \cdot x \leqslant b_i$ is not redundant, we can also pick x_2 with

$$A^{=}x_2 = b^{=} \text{ and } A'x_2 < b',\ a_i \cdot x_2 > b_i$$

Now we can pick $x_0 := tx_1 + (1-t)x_2$ for some $t \in (0,1)$ such that $a_i \cdot x_0 = b_i$.

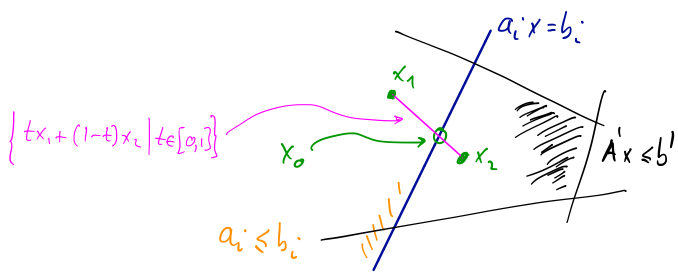

□

Proposition 5.3.12. *Every face of a polyhedron* P *is an intersection of facets of* P.

Proof. Let F be a face of P. By definition, it can be written as $F = \{x \in P \mid A'x = b'\}$ for some subsystem $A' \leqslant b'$ of $Ax \leqslant b$, say consisting of the inequalities $a_i' \cdot x \leqslant b_i'$ for $i = 1, \dots, k$. By Remark 5.3.9 we can assume that all $a_i' \cdot x \leqslant b'$ are effective, irredundant inequalities. Then, obviously

$$F = \bigcap_{i=1}^{k} \{x \in P \mid a_i' \cdot x \leqslant b_i'\}$$

and Lemma 5.3.11 says that every set on the right-hand side is a facet of P. □

5.3.3 Vertices

Lemma 5.3.13. *A polyhedron* P *is an affine subspace if and only if it has no nonempty faces except* P *itself.*

Proof. If P is an affine subspace then it has the form $P = \{Mx = v\}$ for some matrix M and some vector v. Its standard description via inequalities is $P = \{Ax \leqslant b\}$ with $A = \begin{bmatrix} M \\ -M \end{bmatrix}$ and $b = \begin{bmatrix} v \\ -v \end{bmatrix}$. Now obviously the set P satisfies $A'x = b'$ for all subsystems of $Ax \leqslant b$, so every nonempty face equals P.

On the other hand, if $P = \{Ax \leqslant b\}$ has no nonempty faces except P, in particular for the "full" subsystem we have

$$P = \{x \in P \mid Ax = b\} = \{x \mid Ax \leqslant b,\ Ax = b\} = \{Ax = b\},$$

an affine subspace. □

Corollary 5.3.14. *The minimal nonempty faces of a polyhedron are affine subspaces.*

Proof. From Definition 5.3.1 one gathers that every face F of P is a polyhedron in its own right, and every face $F' \subseteq F$ is a face of F if and only if it is a face of P. Thus the minimal nonempty faces of P are polyhedra that do not have nontrivial faces themselves, and by Lemma 5.3.13 are affine subspaces. □

Corollary 5.3.15. *The minimal nonempty faces of a convex polytope* P *have dimension* 0 *(and are called "vertices" of* P*). If* $P = \{Ax \leqslant b\}$*, every vertex of* P *is the unique solution of some system* $A'x = b'$*, where the rows of* A' *are a maximal linearly independent set of the rows of* A*.*

Proof. Since every convex polytope is bounded, the only affine subspaces it can contain are 0-dimensional ones. □

5.4 Unimodularity and integrality

Definition 5.4.1. Call a polytope $P \subseteq \mathbb{R}^n$ *integral* if all its vertices are points in $\mathbb{Z}^n$, *rational* if its vertices lie in $\mathbb{Q}^n$.

We are interested in finding conditions for a polytope to be integral - in particular, we are aiming at proving that Equation (5.2) defines the interior of an integral polytope. To this end, we introduce the notion of unimodularity of matrices.

Let A be a matrix. A *minor* of A is any matrix obtained by deleting some rows and some columns from A.

Definition 5.4.2. An integral matrix A is called *totally unimodular* if every square minor of A has determinant 0, 1 or -1.

Proposition 5.4.3 (Franklin and Veblen, 1921)**.** *Let* A *be a matrix with entries from* $\{0, +1, -1\}$ *and such that each column contains exactly one* $+1$ *and exactly one* -1*. Then* A *is totally unimodular.*

Proof. Let M be a square minor of A. If M has a column with at most one nonzero entry, developing its determinant with respect to that column proves the claim for M. Otherwise every column of M has exactly one $+1$ and one -1: in this case the sum of all rows vanishes, thus M does not have full rank and the determinant of M is 0. □

Corollary 5.4.4. *Incidence matrices of graphs are totally unimodular.*

Proof. See Exercise V.2. □

Our goal in this section is to prove the following characterization of totally unimodular matrices.

Theorem 5.4.5. *For an integral matrix* A*, the following statements are equivalent.*

(i) The matrix A *is totally unimodular*

(ii) For all integral vectors b*, all vertices of the polyhedron* $\{x \mid Ax \leqslant b,\ x \geqslant 0\}$ *are integral.*

(iii) For all integral vectors a, b, c, d*, the polyhedron* $\{x \mid a \leqslant Ax \leqslant b,\ c \leqslant x \leqslant d\}$ *has only integral vertices.*

In order to prove this theorem we start by considering a weaker notion: unimodularity.

Definition 5.4.6. An integer matrix A is called *unimodular* if every maximal square minor has determinant 0, 1 or -1.

Remark 5.4.7. Notice that an $n \times m$-matrix A is totally unimodular if and only if the block matrix $[I|A]$ is unimodular, where I is the $n \times n$-identity matrix (see Exercise V.4)

Remark 5.4.8. Obviously a square matrix B is unimodular if and only if $\det B = \pm 1$. Thus, an invertible square matrix B is unimodular if and only if B^{-1} has only integer entries (as can be ascertained by recalling that $B^{-1} = \frac{1}{\det B} \operatorname{Adj}(B)$ and that the adjugate of an integral matrix is again integral).

Proposition 5.4.9. *Let* A *be an integral matrix of full row rank. Then* A *is unimodular if and only if for every integral vector* b *the polyhedron* $P_b := \{x \mid Ax = b,\ x \geqslant 0\}$ *has only integral vertices.*

Proof.

A *unimodular* $\Rightarrow P_b$ *integral for all integral* b. Suppose that A is a unimodular $n \times m$-matrix and consider a vertex v of P. Then, there are m linearly independent constraints in the system

$$\begin{bmatrix} -I \\ \hline A \\ \hline -A \end{bmatrix} x \leqslant \begin{bmatrix} 0 \\ \hline b \\ \hline -b \end{bmatrix}$$

that are satisfied with equality at v. In particular, the set of columns $\{a_i\}_I$ of A corresponding to nonzero components of v are linearly independent (since, given a linear dependency $\sum_i \lambda_i a_i = 0$, we could add to v a (small enough) multiple of the vector with λ_i in the i-th component for all $i \in I$ and 0 otherwise, and we would obtain a positive-dimensional space W satisfying with equality the same m independent constraints as v, contradicting the fact that v is a vertex). Now since A is full-row-rank we can extend the columns $\{a_i\}_I$ to a column-basis of A, obtaining a maximal (and invertible) square minor B with the property that $B^{-1}b$ equals v at each of its nonzero components. Since A is unimodular, B^{-1} and hence $B^{-1}b$ is integral (see Remark 5.4.8). Since nonzero coordinates of v equal some coordinates of $B^{-1}b$, it follows that v is integral.

P_b *integral for all integral* $b \Rightarrow A$ *unimodular*. Suppose that P_b is integral for all integral b and let B be an invertible maximal minor of A. Consider B^{-1} and, in order to prove $\det B = \pm 1$, we show that $B^{-1}c$ is integral for every integral vector c (see Remark 5.4.8). Take c integral and choose $y \in \mathbb{Z}^n$ such that $z := y + B^{-1}c \geqslant 0$. Then $b := Bz \in \mathbb{Z}^n$ is integral and we can extend z by zero components to a vector $z' \in \mathbb{Z}^m$ with $Az' = Bz = b$. Now z' is a vertex of P_b (as it is in the polyhedron and satisfies n linear independent constraints with equality) and by assumption it is integral. It follows that z (obtained from z' by canceling zero components) and with it $B^{-1}c = z - y$ are integral, as required.

□

Proof. Proof of Theorem 5.4.5

(i) $\Leftrightarrow$ (ii) Recall that A is totally unimodular if and only if $[I|A]$ is unimodular (Remark 5.4.7). By Proposition 5.4.9, this is equivalent to all vertices of the polyhedron $\{z \mid [I|A]z = b,\ z \geqslant 0\}$ being integral, for every integral b. The latter statement is equivalent to saying that all vertices of $\{x \mid x \geqslant 0,\ Ax \leqslant b\}$ (i.e., the projection of the former polyhedron on the last m coordinates) are integral.

(i) ⇔ (iii) Notice that an integer matrix A is unimodular if and only if the matrix

$$\widetilde{A} := \begin{bmatrix} I \\ \hline -I \\ \hline A \\ \hline -A \end{bmatrix}$$

is totally unimodular as well. Now by the equivalence "(i) ⇔ (ii)" this means that

$$\widetilde{P}_f := \{z \mid z \geqslant 0,\ \widetilde{A}x \leqslant f\}$$

has integral vertices for all $f \in \mathbb{Z}^{3n}$. Now given $a, b, c, d \in \mathbb{N}$ let

$$f_0 := \begin{bmatrix} d \\ \hline c \\ \hline b \\ \hline a \end{bmatrix} \tag{5.8}$$

and notice that $\widetilde{P}_{f_0} = \{x \mid a \leqslant Ax \leqslant b,\ c \leqslant x \leqslant d\}$. Since every $f \in \mathbb{Z}^{3n}$ can be written as in Equation (5.8) for suitable a, b, c, d, the proof of equivalence is complete.

□

5.5 Triangulations

We will need to "dissect" a polytope into smaller polytopes.

Definition 5.5.1. A polyhedral (resp. polytopal) complex is a collection $\mathcal{K}$ of polyhedra (resp. polytopes) such that every face of every member of $\mathcal{K}$ is again a member of $\mathcal{K}$, and for any two members $P, Q \in \mathcal{K}$ the intersection $P \cap Q$ is a face of both P and Q. We call faces (resp. vertices, etc.) of $\mathcal{K}$ the faces (vertices, etc.) of members of $\mathcal{K}$. The *support* of the complex is $|\mathcal{K}| := \bigcup_{K \in \mathcal{K}} K$.

We call a set X of points *affinely independent* if $\dim \operatorname{conv} X = |X| - 1$.

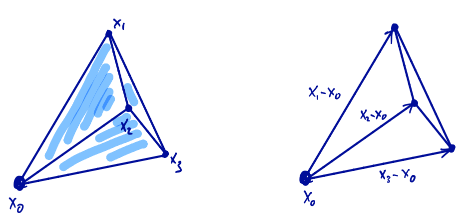

Remark 5.5.2. Notice that given $x_0, \ldots, x_d \in \mathbb{R}^d$, then $\operatorname{conv}\{x_0, \ldots, x_d\} = x_0 + \{\sum_{i=1,\ldots,d} \lambda_i (x_i - x_0) \mid \lambda_i \geqslant 0, \sum_{i>0} \lambda_i \leqslant 1\}$ and this has dimension d if and only if the vectors $x_i - x_0$ are linearly independent in $\mathbb{R}^d$.

In view of the next definition, recall also that a *simplex* is any polytope of the form $\operatorname{conv} X$ for an affinely independent set of points X. In particular, a simplex has the smallest possible number of vertices for a polytope of its dimension.

Definition 5.5.3. Let P be a polytope. We call *subdivision* of P any polytopal complex $\mathcal{K}$ with $|\mathcal{K}| = P$. A *triangulation* of P is a subdivision $\mathcal{T}$ of P such that every member of $\mathcal{T}$ is a simplex.

Our goal is to show that every polytope has a triangulation all of whose vertices are vertices of the polytope. For this, we set up some notation.

Throughout this section let V be a finite subset of $\mathbb{R}^d$, let $P = \operatorname{conv} V$ be a (full-dimensional) polytope in $\mathbb{R}^d$, and let $h : V \to \mathbb{R}_{>0}$ any assignment of a positive real number to every vertex. We think of h as a *height function* and will "lift" the points of V accordingly into $\mathbb{R}^{d+1}$ obtaining a set of points

$$V^{(h)} := \{(v, h(v)) \mid v \in V\} \subset \mathbb{R}^{d+1}.$$

We will throughout write a point in $\mathbb{R}^{d+1} = \mathbb{R}^d \times \mathbb{R}$ as a pair (x, t) with $x \in \mathbb{R}^d$, $t \in \mathbb{R}$. Then, we consider

$$\pi : \mathbb{R}^{d+1} \to \mathbb{R}^d, \quad (x, t) \mapsto x$$

the standard coordinate projection (for instance, $\pi(V^{(h)}) = V$).

Moreover, consider the polyhedron

$$\ell := \{(0, t) \in \mathbb{R}^{d+1} \mid t \geqslant 0\},$$

i.e., the "vertical" coordinate halfline in $\mathbb{R}^{d+1}$ (notice that $\pi(\ell) = 0$).

Definition 5.5.4. The *epigraph* of P with respect to the height function h is

$$P^{\uparrow h} := \operatorname{conv} V^{(h)} + \ell.$$

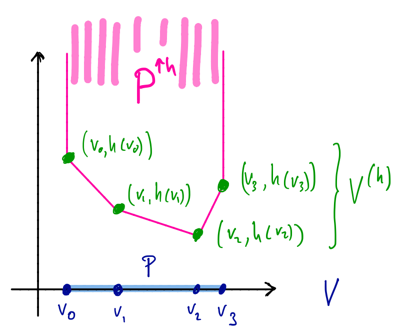

By Proposition 5.2.9, $P^{\uparrow h}$ is a polyhedron. Let us study its supporting hyperplanes.

Lemma 5.5.5. *With the notation above, given a vector* $(a, \alpha) \in \mathbb{R}^{d+1}$ *and* $b \in \mathbb{R}$ *define a hyperplane* $H^{=}$ *and a halfspace* $H^{\leqslant}$ *as*

$$H^{\square} := \{(x, t) \in \mathbb{R}^{d+1} \mid \langle a, x\rangle + \alpha t \;\square\; b\} \qquad \textit{for } \square \in \{=, \leqslant\}.$$

If $P^{\uparrow h} \subseteq H^{\leqslant}$*, then* $\alpha \leqslant 0$*. Moreover, if* $H^{=}$ *is a supporting hyperplane for* $P^{\uparrow h}$*, then the face* $H^{=} \cap P^{\uparrow h}$ *is bounded if and only if* $\alpha < 0$*.*

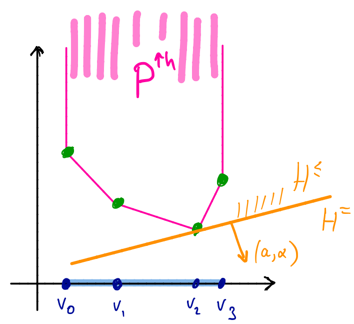

Proof. If $P^{\uparrow h} \subseteq H^{\leqslant}$, then in particular for every $p \in P^{\uparrow h}$ the half-line $p + \ell$ lies in $H^{\leqslant}$, meaning $\langle a, p\rangle + \alpha t \leqslant b$ for all $t \in \mathbb{R}_{+}$. This is only possible if $\alpha \leqslant 0$.

Suppose now that $H^{=}$ is a supporting hyperplane and consider the face $F := H^{=} \cap P^{\uparrow h}$. Now F is unbounded if and only if $F + \ell \subseteq F$. In particular, for every $p \in F$ we have $p + (0, \mathbb{R}_{\geqslant 0}) \subseteq H^{=}$, i.e., $\langle a, p\rangle + \alpha t = b$ for all $t \in \mathbb{R}_{\geqslant 0}$. This is only possible if $\alpha = 0$. □

Corollary 5.5.6. *Every bounded face is a face of some bounded facet.*

Proof. Let F be a bounded face. By Corollary 5.3.15 we know that F is an intersection of facets $F_1, \dots, F_k$: in particular, F is a face of each F_i. Now assume that every F_i is unbounded. By Lemma 5.5.5, this means that F_i is contained in a vertical hyperplane, and with any $x \in F_i$ we have $x + \ell \subseteq F_i$. In particular, also F, the intersection of the F_i, contains the whole halfline $x + \ell$ for every $x \in F$. Since F is non-empty, F would be unbounded: this is a contradiction to all F_i being unbounded. □

Lemma 5.5.7. *It is possible to choose the height function* h *with the following property: if* $\{v_0, \dots, v_d\} \subseteq V$ *is a maximal affinely independent set, then the* $(v_i, h(v_i))$ *are affinely independent and span a hyperplane that does not contain any other point of* $V^{(h)}$*.*

Proof. Write $V = \{v_1, \dots, v_m\}$ and consider any $d+1$-subset $I = \{i_0, \dots, i_d\} \in \binom{[m]}{d+1}$ such that $\{v_i\}_{i \in I}$ is affinely independent. In order for $\{(v_i, h_i)\}_{i \in I}$ to determine a hyperplane in $\mathbb{R}^{d+1}$ they must be affinely independent, and thus (by Remark 5.5.2) we must have that the determinant

$$\delta(I) := \begin{vmatrix} \nu_{i_1} - \nu_{i_0} & \cdots & \nu_{i_d} - \nu_{i_0} & 0 \\ h_{i_1} - h_{i_0} & \cdots & h_{i_d} - h_{i_0} & 1 \end{vmatrix}$$

is non-zero (notice that the last column is never a nontrivial linear combination of the first d, since $\nu_{i_1} - \nu_{i_0}, \dots, \nu_{i_d} - \nu_{i_0}$ are linearly independent by assumption). This is however automatically ensured since this determinant equals $\pm|\nu_{i_1} - \nu_{i_0} \cdots \nu_{i_d} - \nu_{i_0}| \neq 0$ because of affine independence of the ν_i.

Now, if some (ν_j, h_j), $j \notin I$, lies on the same hyperplane as $(\nu_{i_0}, h_{i_0}), \dots, (\nu_{i_d}, h_{i_d})$, then

$$\gamma(I \cup \{j\}) := \begin{vmatrix} \nu_{i_0} & \cdots & \nu_{i_d} & \nu_j \\ h_{i_0} & \cdots & h_{i_d} & h_j \\ 1 & \dots & 1 & 1 \end{vmatrix} = 0$$

Thus, in order to satisfy the claim, we have to choose $h \in \mathbb{R}^m$ so that

$$\prod_{I,\, j \notin I} \gamma(I \cup \{j\}) \neq 0$$

where the product is over all index sets I of maximal affinely independent points in V.

As the expression on the left-hand side is a non-constant polynomial, the set of "good" h is nonempty. ☐

Theorem 5.5.8. *Every polyhedron* $P = \operatorname{conv}(V)$ *has a triangulation all of whose vertices are in* V.

Proof. Choose h as in Lemma 5.5.7 and consider the associated $P^{\uparrow h}$. Define

$$\mathcal{T} := \left\{ \pi(F) \mid F \text{ is a bounded face of } P^{\uparrow h} \right\}$$

Clearly the 0-dimensional members of $\mathcal{T}$ are a subset of V. Moreover, for every bounded facet F of $P^{\uparrow h}$, say supported on a hyperplane H (say $F = H^{=} \cap P^{\uparrow h}$ and $P^{\uparrow h} \subseteq H^{\leqslant} = \{\langle (a, \alpha), (x, t) \rangle \leqslant b\}$):

(a) F is a simplex.
Proof. By the choice of h we know that every bounded facet F of $P^{\uparrow h}$ is supported on a hyperplane that contains at least $d+1$ affinely independent vertices (since $\dim(F) = d$), but no hyperplane can contain more than $d+1$ points from $V^{(h)}$: thus F is a simplex.

(b) The restriction $\pi : F \to \pi(F)$ is bijective, with inverse $\iota : x \mapsto \left(x, \frac{b - \langle a, x \rangle}{\alpha}\right)$.
Proof. Since F is bounded, by Lemma 5.5.5 $\alpha \neq 0$. Thus the inverse is well-defined and proves bijectivity.

(c) For every $X \subseteq F$, $\mathrm{conv}(\pi(X)) = \pi(\mathrm{conv}(X))$

Proof. Given $(q_0, t_0), \ldots, (q_k, t_k) \in F$ and $\lambda_1, \ldots, \lambda_k \in \mathbb{R}$, we have $(q_0, t_0) = \sum_i \lambda_i (q_i, t_i)$ if and only if $q_0 = \sum_i \lambda_i q_i$ (for the "if" direction: recall the inverse map of item (b) and notice that $t_i = (b - \langle a, q_i \rangle)/\alpha$ for all i, and thus $\alpha \sum_{i \geqslant 1} \lambda_i t_i = b - \langle a, \sum_{i \geqslant 1} \lambda_i q_i \rangle = b - \langle a, q_0 \rangle = \alpha t_0$)

(d) $\pi(F)$ is a simplex with set of vertices $\pi(V_F)$, where V_F is the set of vertices of F.

Proof. Since F is a simplex of dimension d, $|V_F| = d + 1$. Then, by item (c) $\pi(F) = \mathrm{conv}\, \pi(V_F)$ has at most $d+1$ vertices. Moreover, linearity of π implies that $\dim \pi(F) = \dim(F) = d$ and, in particular, $\pi(F)$ has at least $d+1$ vertices. The claim follows.

(e) The faces of $\pi(F)$ correspond bijectively to projections of faces of F.

Proof. Faces of simplices correspond bijectively to (convex hulls of) subsets of vertices. Per items (a) and (d) both F and $\pi(F)$ are simplices with vertices, respectively, V_F and $\pi(V_F)$. Thus the bijection $2^{V_F} \to 2^{\pi(V_F)}$, $I \mapsto \pi(I)$, induces a bijection between the faces of F and $\pi(F)$.

Moreover,

(f) $\mathcal{T}$ consists of all faces of projections of bounded facets.

Proof. By Corollary 5.5.6 every bounded face of $P^{\uparrow h}$ is a face of a bounded facet. Thus, by item (e), $\mathcal{T}$ is the set of all faces of some $\pi(F)$, F a bounded facet of F.

We now claim that $\mathcal{T}$ is a triangulation of P. For this we need to check three things.

(1) $|\mathcal{T}| = P$.
The inclusion $\subseteq$ follows trivially from $P = \pi(P^{\uparrow h})$.

In order to prove the inclusion $\supseteq$ is is enough to prove that the open interior of P is contained in $\pi(P^{\uparrow h})$, as π is continuous. Let then p be a point in the interior of P and let ϵ be such that the ϵ-ball $B_\epsilon(p)$ is fully contained in P. Moreover, let

$$\ell_p := (p + \ell) \cap P^{\uparrow h}, \text{i.e.,} \ell_p = x + \ell$$

for some $x \in P^{\uparrow h}$. Then x is on the boundary of $P^{\uparrow h}$ and therefore it is in the intersection of $P^{\uparrow h}$ with a supporting hyperplane. In the notation of Lemma 5.5.5, there is a vector (a, α) (which we can suppose to be of norm 1) and a number b such that $H^=$ is a supporting hyperplane of $P^{\uparrow h}$ and

$$x \in F := P^{\uparrow h} \cap H^{\leqslant}.$$

Now, if α=0 then $\pi(H^{=})$ is a hyperplane in $\mathbb{R}^d$ that contains p and bounds the halfspace $\pi(H^{\leqslant})$. Now $P = \pi(P^{\uparrow h}) \subseteq \pi(H^{\leqslant})$, but $p + \frac{\epsilon}{2}a \notin \pi(H^{\leqslant})$, a contradiction to $B_\epsilon(p) \subseteq P$.

This means that $\alpha > 0$ and, by Lemma 5.5.5, the face F is bounded. Therefore $p = \pi(x) \in \pi(F) \subseteq |T|$ as was to be shown.

(2) Every member of $\mathcal{T}$ is a simplex.

Item (f) shows that every member of $\mathcal{T}$ is a face of $\pi(F)$ for some bounded facet F. Since every face of a simplex is a simplex, item (d) ensures that every member of $\mathcal{T}$ is a simplex.

(3) $\mathcal{T}$ is a polytopal complex.

Since every face of a bounded face is bounded, item (f) shows that $\mathcal{T}$ contains all faces of each of its members. We are left to show that the intersection of any two $T_1, T_2 \in \mathcal{T}$ is a face of both. Let F_1 and F_2 be bounded faces of $P^{\uparrow h}$ such that $T_i = \pi(F_i)$ for $i = 1, 2$. By item (e) it is enough to prove

$$\pi(F_1 \cap F_2) = T_1 \cap T_2.$$

The inclusion $\subseteq$ is trivial.

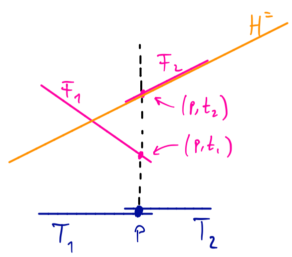

Furthermore, for any $p \in T_1 \cap T_2 \setminus \pi(F_1 \cap F_2)$ the halfline $p + \ell \supseteq$ intersects both F_is. Therefore there are t_1, t_2 such that $(p, t_1) \in F_1$ and $(p, t_2) \in F_2$, and $t_1 \neq t_2$ because p is not in $\pi(F_1 \cap F_2)$. Without loss of generality we can assume $t_1 < t_2$. Consider a supporting hyperplane $H = \{(a, \alpha) \cdot x = \delta\}$ to F_2. Then $(p, t_2) \in F_2$ implies $\langle a, p \rangle + \alpha t_2 = \delta$. Since F_2 is bounded, we have $\alpha < 0$ and so $\langle a, p \rangle + \alpha t_1 > \delta$, contradicting the fact that, since H is a supporting hyperplane, we should have $P \subseteq H^{\leqslant}$. Thus $T_1 \cap T_2 \subseteq \pi(F_1 \cap F_2)$ and we are done.

□

Definition 5.5.9. We say that a polyhedron P (resp. a pointed cone) *can be triangulated using no new vertices* if there exists a triangulation $\mathcal{T}$ of P such that the vertices of $\mathcal{T}$ are vertices of P.

Analogously, we say that a pointed cone C *can be triangulated using no new generators* if there exists a triangulation $\mathcal{T}$ of C such that the generators of $\mathcal{T}$ are generators of C.

Corollary 5.5.10. *Every polyhedron P can be triangulated using no new vertices.*

5.6 Sources and Bibliography

The approach of Section 5.1 is from the original paper by Kochol [5]. Sections 5.2 to 5.4 follow more or less closely the treatment by Schrijver [6]. Section 5.5 is based on Chapter 5 of Beck and Sanyal [2] and Chapter 3 of Beck and Robins [1]

Exercises V

V.1 This exercise should give a geometric intuition for the "Farkas' type" statements we have seen. Let

(a) Consider the statement of Farkas' lemma. Let

$$A := \begin{bmatrix} 1 & 2 \\ -1 & -1 \end{bmatrix}, \qquad b_1 := \begin{bmatrix} 8 \\ -2 \end{bmatrix}, \qquad b_2 := \begin{bmatrix} 8 \\ -5 \end{bmatrix}.$$

For $i = 1, 2$ do the following: determine the solutions of $Ax = b_i$; then draw the set of points y with $yA \geqslant 0$ and check whether $yb_i \geqslant 0$ for all such y.

(b) In order to illustrate Corollary 5.3.5, consider

$$A := \begin{bmatrix} 1 & 1 \\ 1 & -2 \\ -2 & 1 \end{bmatrix}, \qquad b_1 := \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \qquad b_2 := \begin{bmatrix} -1 \\ 0 \\ 0 \end{bmatrix}.$$

For $i = 1, 2$ determine the set of solutions of the system $Ax \leqslant b_i$, consider the set of all $y \in \mathbb{R}^3$ with $y \geqslant 0$ and $yA = 0$ and verify whether $yb \geqslant 0$ for all such y.

(c) Check the claim of Proposition 5.3.6 using the data A from (b) above and $b := \begin{bmatrix} 1 \\ 2 \\ 0 \end{bmatrix}$, $c := \begin{bmatrix} 1 \\ 1 \end{bmatrix}$. In particular, draw a picture of $Ax \leqslant b$ and of the points where $c \cdot x$ is maximum; then draw the set $yA = c$, $y \geqslant 0$ and compute the minimum on the right-hand side.

V.2 Recall the definition of the incidence matrix of a graph and prove Corollary 5.4.4.

V.3 (+) Let P be a polyhedron and consider the set $\mathscr{F}(P)$ of all faces of P, partially ordered by inclusion.

(a) Let $x_0, \ldots, x_d \in \mathbb{R}^d$ be chosen so that these points do not lie all on the same (affine) hyperplane. Then $P := \operatorname{conv}\{x_0 \, \mathtt{ldots}, x_d\}$ is called a *simplex*. What is the poset of faces of a simplex? (You can start by concrete examples with $d = 2, 3$).

(b) Prove that for every polyhedron P, the poset $\mathscr{F}(P)$ is ranked.

(c) $^{++}$ Prove that for every convex polytope P the poset $\mathscr{F}(P)$ is a ranked lattice.

V.4 Prove that an integer $n \times m$ matrix A is totally unimodular if and only if the block matrix $[1 \mid A]$ is unimodular (here 1 denotes the $n \times n$ identity matrix).

V.5 ([1, Exercise 3.3]) Let $e_1, \dots, e_d$ be the standard basis vectors of $\mathbb{R}^d$. For every permutation π of the set $[n]$ consider the polytope

$$\Delta_\pi := \operatorname{conv}\{0, e_{\pi(1)}, e_{\pi(1)} + e_{\pi(2)}, \dots, e_{\pi(1)} + e_{\pi(2)} + \dots + e_{\pi(d)}\}.$$

Prove that the polytopes Δ_π and all their faces define a triangulation of the unit cube $[0,1]^d$. Moreover, prove that any two such simplices can be obtained from one another by a congruence (a concatenation of reflections, rotations and translations).

Chapter 6

Rational generating functions

The last ingredient we need towards Ehrhart's theorem is rational generating functions as "representations" of number sequences.

6.1 Basics

Fix a natural number d and a field $\mathbb{K}$.

The main idea is that any function $f : \mathbb{Z}^d \to \mathbb{K}$ can be represented by the formal expression in the variables $z_1, \ldots, z_d$:

$$\sum_{m \in \mathbb{Z}^d} f(m) z^m, \qquad \text{where we write } z^m := z_1^{m_1} \cdots z_d^{m_d}.$$

Every such expression is called *formal Laurent series*. The word "formal" indicates that we are not concerned about analytic convergence. We treat such "sum" as abstract formal objects.

It is common to call

$$\mathbb{C}[[z_1, \ldots, z_d]] := \left\{ \sum_{m \in \mathbb{N}} f(m) z^m \,\middle|\, f : \mathbb{N}^d \to \mathbb{C} \right\},$$

the set of *formal power series* in the variables $z_1, \ldots, z_d$. The *degree* of $\sum_{m \in \mathbb{Z}^d} f(m) z^m$ is defined as $\max\{\sum_i m_i \mid f(m) \neq 0\} \in \mathbb{N} \cup \{\infty\}$. Thus an element of $\mathbb{C}[[z_1, \ldots, z_d]]$ of finite degree is a polynomial, i.e., an element of $\mathbb{C}[z_1, \ldots, z_d]$.

6.1.1 Algebraic structure

The following addition and multiplication rules define a ring structure on the set of formal power series (and of formal Laurent series, by changing the index set from $\mathbb{N}$ to $\mathbb{Z}$):

$$\sum_{m\in\mathbb{N}} f(m)z^m + \sum_{m\in\mathbb{N}} g(m)z^m := \sum_{m\in\mathbb{N}} (f(m)+g(m))z^m$$

$$\sum_{m\in\mathbb{N}} f(m)z^m \cdot \sum_{m\in\mathbb{N}} g(m)z^m := \sum_{m\in\mathbb{N}} \left(\sum_{k\in\mathbb{Z}} f(k)g(m-k)\right) z^m.$$

Example 6.1.1. The following identity holds in the ring $\mathbb{C}[[z]]$:

$$\frac{1}{1-z} = \sum_{k\in\mathbb{N}} z^k.$$

Thus, the function $f : \mathbb{N} \to \mathbb{C}$ constant equal to 1 can be represented by the fraction on the left-hand side. We call "rational" any formal power series that equals a rational expression in the ring of formal power series.

The gist now is: *whether a function $f : \mathbb{N} \to \mathbb{C}$ agrees with a polynomial can be ascertained by looking at the rationality of the associated formal power series.* In the next section we will make this precise.

Before that, let us collect some considerations on the case $\mathbb{K} = \mathbb{C}$

6.2 The main theorem

Theorem 6.2.1. *Fix $\alpha_1, \ldots, \alpha_d \in \mathbb{C}$ with $d \geqslant 1$ and $\alpha_d \neq 0$. Then, for every function $f : \mathbb{N} \to \mathbb{C}$ the following are equivalent:*

(i)

$$\sum_{n\in\mathbb{N}} f(n)z^n = \frac{p(z)}{q(z)}$$

where $q(z) = 1 + \alpha_1 z + \ldots + \alpha_d z^d$ and $p(z)$ is any polynomial in z of degree less than d.

(ii) For all $n \in \mathbb{N}$:

$$\sum_{i=0}^{d} \alpha_i f(n+d-i) = 0$$

(iii) For all $n \in \mathbb{N}$,

$$f(n) = \sum_{i=1}^{k} p_i(n)\gamma_i^n$$

where the γ_i are distinct, nonzero complex numbers with $1 + \alpha_1 z + \ldots + \alpha_d z^d = (1-\gamma_1 z)^{d_1} \cdots (1-\gamma_k z)^{d_k}$ and each p_i is a polynomial of degree less than d_i

Proof. Fix $q(z) := 1 + \alpha_i z + \ldots + \alpha_d z^d = (1-\gamma_1 z)^{d_1} \cdots (1-\gamma_k z)^{d_k}$, all γ_i nonzero and distinct. Consider the following complex vectorspaces:

- $V_1 := \{f : \mathbb{N} \to \mathbb{C}$ such that (i) holds$\}$.
  Notice that $\dim V_1 = d$, as in (i) we can choose the d coefficients of $p(z)$ freely.

- $V_2 := \{f : \mathbb{N} \to \mathbb{C}$ such that (ii) holds$\}$.
  Here we have $\dim V_2 = d$ because in (ii) the choice of the $f(n+i)$, $i = 0, \ldots, d-1$ is free and determines f completely.

- $V_3 := \{f : \mathbb{N} \to \mathbb{C}$ such that (iii) holds$\}$.
  Here, for every i we can choose the d_i coefficients of p_i freely, and this choice determines f completely. Thus, $\dim V_3 = d_1 + \ldots + d_k = d$.

- $$V_4 := \left\{ \begin{array}{ll} f : \mathbb{N} \to \mathbb{C} \text{ such that} & \\ \sum_{n\in\mathbb{N}} f(n) z^n = \sum_{i=1}^{k} \sum_{j=1}^{d_i} \beta_{ij} (1-\gamma_i z)^{-j} & (\sharp) \\ \text{for some } \beta_{ij} \in \mathbb{C} \text{ and where the } \gamma_i, d_i \text{ are as in (iii)} & \end{array} \right\}.$$
  We claim that $\dim V_4 = d$. This is not as evident as in the other cases.

  First, notice that V_4 is spanned over $\mathbb{C}$ by the rational functions

  $$R_{ij}(z) := \frac{1}{(1-\gamma_i z)^j}, \qquad i = 1, \ldots, k,\ j = 1, \ldots, d_i.$$

  Since there are $\sum_i d_i = d$ such R_{ij}s, we have $\dim V_4 \leqslant d$.

  Now it is enough to show that the R_{ij} are linearly independent. By way of contradiction suppose that there is a linear dependency $\sum c_{ij} R_{ij}(z) = 0$ with complex coefficients c_{ij} not all equal to zero. Let i_0 such that $c_{i_0 j} \neq 0$ for some j, and let j_0 be the largest such index. Now multiplying the linear dependency by $(1-\gamma_{i_0} z)^{j_0}$ and setting $z = \frac{1}{\gamma_{i_0}}$ ([1]) gives the equality $c_{i_0 j_0} = 0$, a contradiction. Thus the R_{ij} are linearly independent and $\dim V_4 = d$.

Now it is enough to show that $V_1 = V_2 = V_3 = V_4$. We now do so, in three steps.

$V_3 = V_4$ Since $\dim V_3 = \dim V_4$ it is enough to prove $V_4 \subseteq V_3$. To this end, recall that

$$\frac{1}{(1-\gamma_i z)^j} = \left(\sum_{n\in\mathbb{N}} \gamma_i^n z^n \right)^j = \sum_{n\in\mathbb{N}} \binom{j+n-1}{j-1} \gamma_i^n z^n.$$

[1]Formal power series are formal objects, and one has to make sense of "setting z to...". In this case, the subring $\mathcal{R}$ of $\mathbb{C}[[z]]$ generated by the R_{ij} with $i \neq i_0$ and $(1-\gamma_{i_0} z)$ is in the image of the ring of rational functions $\mathbb{C}(z)$, and the map $\mathbb{C}(z) \to \mathbb{C}[[z]]$ is injective. Now every function in $\mathcal{R}$ is defined for $\frac{1}{\gamma_{i_0}}$ and thus the map $\mathcal{R} \to \mathbb{C}$ obtained by formally replacing z with $\frac{1}{\gamma_{i_0}}$ is well-defined, so that an identity in $\mathcal{R}$ goes to an identity in $\mathbb{C}$.

Notice that $\binom{j+n-1}{j-1}$, as a function of n, is a polynomial of degree $j-1$. Now let $f \in V_4$. For all n, $f(n)$ equals the coefficient of z^n in ($\sharp$), which we can compute to be

$$\sum_{i=1}^{k} \left(\sum_{j=1}^{d_i} \beta_{ij} \binom{j+n-1}{j-1} \right) \gamma_i^n, \qquad (\flat)$$

where the expression inside the big parenthesis is a polynomial in n of degree at most $d_i - 1$. Therefore $f \in V_3$ and the proof is complete.

$V_4 = V_1$ The right-hand side of ($\sharp$) equals

$$\frac{\sum_{i=1}^{k} \left(\prod_{h \neq i} (1-\gamma_h z)^{d_h} \sum_{j=1}^{d_i} \beta_{ij} (1-\gamma_i z)^{d_i - j} \right)}{\prod_{i=1}^{k} (1-\gamma_i z)^{d_i}}, \qquad (\flat\flat)$$

where the numerator is a polynomial of degree strictly smaller than $\sum_i d_i = d$. Thus, $V_4 \subseteq V_1$. Since both spaces have dimension d, equality follows.

$V_1 = V_2$ If $f \in V_1$, then

$$q(z) \sum_{n \in \mathbb{N}} f(n) z^n = p(z).$$

Now equating the coefficient of z^d on both sides of the last equality gives the relation in (ii).

□

Call a function $f : \mathbb{N} \to \mathbb{C}$ *polynomial* if there is a polynomial $p \in \mathbb{C}[z]$ with $f(n) = p(n)$ for all $n \in \mathbb{N}$.

Corollary 6.2.2. *Let* $f : \mathbb{N} \to \mathbb{C}$ *and* $d \in \mathbb{N}$. *Then* f *is polynomial of degree at most* d *if and only if*

$$\sum_{n \in \mathbb{N}} f(n) z^n = \frac{p(z)}{(1-z)^{d+1}} \qquad (\natural)$$

for some $p(x) \in \mathbb{C}[x]$ *of degree at most* d.

Moreover, f *is polynomial of degree exactly* d *if and only if* $p(1) \neq 0$.

Proof. Equation ($\natural$) is an instance of Theorem 6.2.1.(i), with $k = 1$ and $\gamma_1 = 1$. From the equivalence of (i) and (iii) in Theorem 6.2.1 we have then that ($\natural$) is equivalent to

$$f(n) = p_1(n)$$

for a polynomial $p_1 \in \mathbb{C}[x]$ of degree at most d.

For the second claim notice that following $(\flat)$ in the proof of Theorem 6.2.1 our f satisfies

$$f(n) = \sum_{j=1}^{d+1} \beta_{1j} \binom{j+n-1}{j-1}$$

for some numbers β_{1j}, so f is polynomial of degree d if and only if $\beta_{1(d+1)} \neq 0$. Now, the numbers β_{ij} are related to $p(z)$ via the expression in $(\flat\flat)$. In our case, this means

$$p(z) = \sum_{j=1}^{d+1} \beta_{ij} (1-z)^{d+1-j}$$

and from here $p(1) \neq 1$ if and only if $\beta_{1(d+1)} \neq 0$. The claim follows. □

Exercises VI

VI.1 Give a nonrecursive form for the general term of the following recursively defined integer sequences.

- $a_0 = 2$, $a_1 = 3$, $a_n = 3a_{n-1} - 2a_{n-2}$ for $n \geqslant 2$
- $a_0 = 0$, $a_1 = 2$, $a_n = 4a_{n-1} - 4a_{n-2}$ for $n \geqslant 2$
- $a_0 = 5$, $a_1 = 12$, $a_n = 4a_{n-1} - 3a_{n-2} - 2^{n-2}$ for $n \geqslant 2$

Chapter 7

Ehrhart theory

NOTE: This chapter follows very closely chapter 3 of Beck-Robins “Computing the continuum discretely". I omit proofs and examples that can be found in that chapter.

Our goal in this chapter is to determine the *lattice-point enumerator* function of a given convex polytope P in $\mathbb{R}^d$, defined for every *integer* t as

$$L_P(t) := \#(tP \cap \mathbb{Z}^d) = \#(P \cap \frac{1}{t}\mathbb{Z}^d),$$

i.e., as the number of integer points contained in the t-th dilate of P.

This chapter is devoted to proving

Theorem 7.0.1. *For every integral* d*-dimensional polytope* P*,* $L_P(t)$ *is polynomial of degree* d*.*

The preceding chapter should suggest to us to consider the *Ehrhart series* in one variable z:

$$\mathrm{Ehr}_P(z) := \sum_{t \geqslant 0} L_P(t) z^t$$

7.1 Formal Laurent series from integer points

Definition 7.1.1. Let P be any polyhedron in $\mathbb{R}^d$. Define

$$\sigma_P(z_1, \ldots, z_d) := \sum_{m \in P \cap \mathbb{Z}^d} z^m.$$

We call this the *integer-point transform* of P.

Example 7.1.2. See Examples 3.3 and 3.4 in the book by Beck and Robins [1]

Recall that a *simplicial* cone is a polyhedral cone C generated by $\dim(C)$ vectors.

Theorem 7.1.3. *Let* $m_1, \ldots, m_d \in \mathbb{Z}^d$ *be such that*

$$C := \text{cone}\{m_1, \ldots, m_d\}$$

is a simplicial d*-dimensional cone. Let*

$$\Pi := \left\{ \sum_{i=1}^{d} \lambda_i m_i \mid 0 \leqslant \lambda_i < 1 \right\}$$

be the "fundamental parallelepiped" of C. *Then, for every* $v \in \mathbb{R}^n$ *we have*

$$\sigma_{v+C}(z_1, \ldots, z_d) = \frac{\sigma_{v+\Pi}(z_1, \ldots, z_d)}{(1 - z^{m_1}) \cdots (1 - z^{m_d})}$$

Proof. The proof follows the template of the previous example - see [1]. □

7.2 Coning over polytopes

Definition 7.2.1. Let $P \subseteq \mathbb{R}^d$ be a convex polytope with vertices $x_1, \ldots, x_k$. We call *standard cone over* P the polyhedron

$$C(P) := \text{cone}\{(x_1, 1), \ldots (x_k, 1)\} \subseteq \mathbb{R}^{d+1}.$$

Now consider the ("vertical") vector $e_{d+1} := (0, \ldots, 0, 1) \in \mathbb{R}^{d+1}$. Given an integer t we let

$$H^{(t)} := \{y \in \mathbb{R}^{d+1} \mid \langle e_{d+1}, y \rangle = t\}$$

denote the "horizontal" plane at "height" t and define

$$P^{(t)} := C(P) \cap H^{(t)}.$$

Remark 7.2.2. Notice that, under the canonical identification of $H^{(t)}$ with $\mathbb{R}^d$, we have $P^{(1)} = P$ and, in general, $P^{(t)} = tP$, the t-th dilate of P. Thus, for the integer point counting function we have

$$L_P(t) = L_{P^{(t)}}(1)$$

and, since every integer point of $C(P)$ is obviously an integer point of exactly one $P^{(t)}$,

$$\sigma_{C(P)}(1, \ldots, 1, z) = 1 + \sum_{t \geqslant 1} L_{P^{(t)}}(1) z^t = \sum_{t \geqslant 0} L_P(t) z^t = \text{Ehr}_P(z)$$

Proof of Theorem 7.0.1

By Corollary 6.2.2 it is enough to prove that the power series satisfies

$$\mathrm{Ehr}_P(z) = \frac{p(z)}{(1-z)^{d+1}}$$

for a polynomial p of degree at most d and $p(1) \neq 0$.

Moreover, by Theorem 5.5.8 it is enough to prove this for P an integral simplex (since any integral P can be triangulated with no new vertices, and in particular the triangulation consists of integral simplices).

Then, let P be an integral simplex – say with vertices $x_1, \ldots, x_{d+1} \in \mathbb{R}^d$. Then $C(P)$ is a simplicial d-cone generated by

$$m_1 := (x_1, 1), \ldots, m_{d+1} = (x_{d+1}, 1) \in \mathbb{Z}^{d+1}$$

and now from Remark 7.2.2 and Theorem 7.1.3 that

$$\mathrm{Ehr}_P(z) = \sigma_{C(P)}(1, \ldots 1, z_{d+1}) = \frac{\sigma_\Pi(1, \ldots, 1, z_{d+1})}{(1-z)^{d+1}}$$

where $\Pi = \{\lambda_1 m_1 + \ldots + \lambda_{d+1} m_{d+1} \mid 0 \leqslant \lambda_i < 1\}$ is the half-open fundamental parallepiped.

Now $\sigma_\Pi(1, \ldots 1, z_{d+1})$ is a polynomial in z_{d+1} (since Π is bounded), and its value at $z_{d+1} = 1$ is not zero (in fact, $\sigma_\Pi(1, \ldots 1)$ is the total number of integer points in Π, and Π contains at least the origin). We only have to show that the degree of $\sigma_\Pi(1, \ldots 1, z_{d+1})$ is at most d. For this, notice that the maximum z_{d+1}-degree in $\sigma_\Pi(z_1, \ldots, z_{d+1})$ is the maximum value of the last coordinate over all integer points in Π. Now the last coordinate of a generic point $\lambda_1 m_1 + \ldots + \lambda_{d+1} m_{d+1}$ in the parallelepiped is $\lambda_1 + \ldots + \lambda_{d+1}$. Thus, every point in Π has last coordinate strictly less than $d+1$, and so every integer point in Π has last coordinate at most d.

Exercises VII

VII.1 For each of the following simplices compute the Ehrhart series and the Ehrhart polynomial.

- $\text{conv}\{(0,0),(1,0),(0,1)\}$
- $\text{conv}\{(0,0,0),(1,0,0),(0,2,0),(0,0,3)\}$
- $\text{conv}\{(0,0,0,0),(1,0,0,0),(0,2,0,0),(0,0,3,0),(0,0,0,4)\}$

VII.2 Compute the Ehrhart polynomial of the unit square $[0,1]^2$ in $\mathbb{R}^2$ using the theory developed in this chapter, and compare the result with the "common sense" computation.

VII.3 Let G be the graph depicted below, oriented so that "every edge is oriented clockwise"

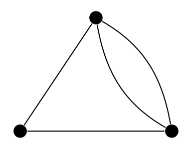

Compute the number of positive k-flows on G using Ehrhart theory, following the steps below:

- Consider the incidence matrix A of G and compute the vertices of the polytope $P = \{Ax = 0,\ \mathbf{0} \leqslant x \leqslant \mathbf{1}\}$ (this is the polytope $\overline{U}$ of the beginning of chapter V).

  (Hint: this polytope lives in the 2-dimensional space $W = \ker(A) \subseteq \mathbb{R}^4$, hence the candidates for being vertices are the intersections of the lines $W \cap \{x_i = 0\}$, $W \cap \{x_i = 1\}$, $i = 1, \ldots, 4$.)
- Compute the integer point enumerator $L_Q(t)$ for Q ranging over all faces of P.
- Give a polynomial expression for the function $I_G(k)$, counting the number of integer points in U.

Bonus question: how many flippings of G support positive k-flows? Can you compute the "full" integer flow polynomial of G?

Bibliography

[1] Matthias Beck and Sinai Robins. *Computing the continuous discretely*. Undergraduate Texts in Mathematics. Springer, New York, second edition, 2015. Integer-point enumeration in polyhedra, With illustrations by David Austin.

[2] Matthias Beck and Raman Sanyal. *Combinatorial reciprocity theorems*, volume 195 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2018. An invitation to enumerative geometric combinatorics.

[3] Reinhard Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, Berlin, fifth edition, 2018. Paperback edition of [MR3644391].

[4] Chris Godsil and Gordon Royle. *Algebraic graph theory*, volume 207 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2001.

[5] Martin Kochol. Polynomials associated with nowhere-zero flows. *J. Combin. Theory Ser. B*, 84(2):260–269, 2002.

[6] Alexander Schrijver. *Theory of linear and integer programming*. Wiley-Interscience Series in Discrete Mathematics. John Wiley & Sons, Ltd., Chichester, 1986. A Wiley-Interscience Publication.

[7] W. T. Tutte. *Graph theory*, volume 21 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 2001. With a foreword by Crispin St. J. A. Nash-Williams, Reprint of the 1984 original.

[8] Hassler Whitney. A logical expansion in mathematics. *Bull. Amer. Math. Soc.*, 38(8):572–579, 1932.