Chapter 1

Four Polynomials

To many, mathematics is a collection of theorems. For me, mathematics is a collection of examples; a theorem is a statement about a collection of examples and the purpose of proving theorems is to classify and explain the examples...
John B. Conway

In the spirit of the above quote, this chapter serves as a source of examples and motivation for the theorems to come and the tools to be developed. Each of the following four sections introduces a family of examples together with a combinatorial reciprocity statement which we will prove in later chapters.

1.1. Graph Colorings

Graphs and their colorings are all-time favorites in introductory classes on discrete mathematics, and we too succumb to the temptation to start with one of the most beautiful examples. A **graph** $G = (V, E)$ is a discrete structure composed of a finite[1] set of **nodes** V and a collection $E \subseteq \binom{V}{2}$ of unordered pairs of nodes, called **edges**. More precisely, this defines a **simple** graph as it excludes the existence of multiple edges between nodes and, in particular, edges with equal endpoints, i.e., **loops**. We will, however, need such nonsimple graphs in the sequel but we dread the formal overhead nonsimple graphs entail and will trust your discretion to make the necessary modifications. The most charming feature of graphs is that they are easy to visualize and their natural habitat is the margins of textbooks or notepads. Figure 1.1 shows some examples.

[1] *Infinite* graphs are interesting in their own right; however, they are no fun to color-count and so will play no role in this book.

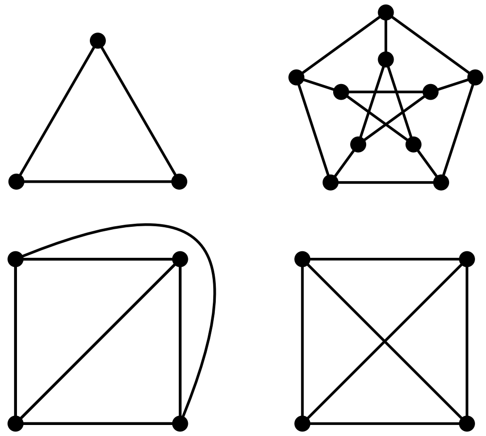

Figure 1.1. Various graphs.

An n-**coloring** of a graph G is a map $c : V \to [n] := \{1, 2, \ldots, n\}$. An n-coloring c is called **proper** if no two nodes sharing an edge get assigned the same color, that is,

$$c(u) \neq c(v) \qquad \text{whenever} \qquad uv \in E\,.$$

The name *coloring* comes from the natural interpretation of thinking of $c(v)$ as one of n possible colors that we use for the node v. A proper coloring is one where adjacent nodes get different colors. Here is a first indication why considering simple graphs often suffices: the existence and even the number of n-colorings is unaffected by parallel edges, and there are simply no proper colorings in the presence of loops.

Much of the fame of graph colorings stems from a question that was asked around 1852 by Francis Guthrie and answered only some 124 years later. In order to state the question in modern terms, we call a graph G **planar** if G can be drawn in the plane (or scribbled in the margin) such that edges do not cross except possibly at nodes. For example, the last row in Figure 1.1 shows a planar and nonplanar embedding of the (planar) graph K_4. Here is Guthrie's famous conjecture, now a theorem.

Four-color Theorem. *Every planar graph has a proper* 4*-coloring.*

There were several attempts at the Four-color Theorem before the first correct proof by Kenneth Appel and Wolfgang Haken. Here is one particularly interesting (but not yet successful) approach to proving the four-color theorem, due to George Birkhoff. For a (not necessarily planar) graph G, let

$$\chi_G(n) \ := \ |\{c : V \to [n] \text{ proper } n\text{-coloring}\}|\,.$$

The following observation, due to George Birkhoff and Hassler Whitney, is that $\chi_G(n)$ is the restriction to $\mathbb{Z}_{>0}$ of a particularly nice function.

Proposition 1.1.1. *If $G = (V, E)$ is a loopless graph, then $\chi_G(n)$ agrees with a polynomial of degree $|V|$ with integral coefficients. If G has a loop, then $\chi_G(n) = 0$.*

By a slight abuse of notation, we identify $\chi_G(n)$ with this polynomial and call it the **chromatic polynomial** of G. Nevertheless, we emphasize that, so far, only the values of $\chi_G(n)$ at positive integral arguments have an interpretation in terms of G.

Birkhoff's motivation to introduce the chromatic polynomial was that the four-color theorem is equivalent to the statement $\chi_G(4) > 0$ for all planar graphs G.

One proof of Proposition 1.1.1 is interesting in its own right, as it exemplifies *deletion–contraction* arguments which we will revisit in Chapter 7. For $e \in E$, the **deletion** of e results in the graph $G \setminus e := (V, E \setminus \{e\})$. The **contraction** G/e is the graph obtained by identifying the two nodes incident to e and removing e. An example is given in Figure 1.2.

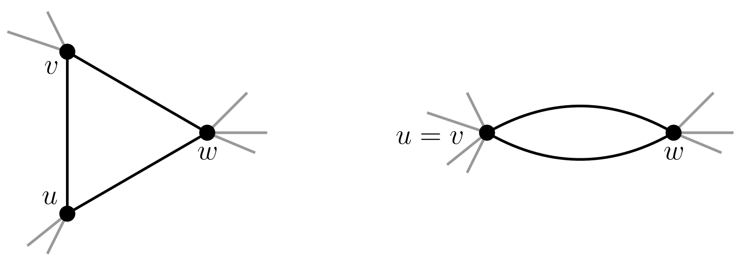

Figure 1.2. Contracting the edge $e = uv$.

Proof of Proposition 1.1.1. If G has a loop, then it admits no proper coloring by definition. For the more interesting case that G is loopless, we induct on $|E|$.

For $|E| = 0$ there are no coloring restrictions and $\chi_G(n) = n^{|V|}$. One step further, assume that G has a single edge $e = uv$. Then we can color all nodes $V \setminus \{u\}$ arbitrarily and, assuming $n \geq 2$, can color u with any color $\neq c(v)$. Thus, the chromatic polynomial is $\chi_G(n) = n^{d-1}(n-1)$, where $d = |V|$.

For the induction step, let $e = uv \in E$. We claim

$$\chi_G(n) = \chi_{G \setminus e}(n) - \chi_{G/e}(n) . \tag{1.1.1}$$

Indeed, a coloring c of $G \setminus e$ fails to be a coloring of G if $c(u) = c(v)$. That is, we are over-counting by all proper colorings that assign the same color to u and v. These are precisely the proper n-colorings of G/e.

By (1.1.1) and the induction hypothesis, $\chi_G(n)$ is the difference of a polynomial of degree $d = |V|$ and a polynomial of degree $\leq d-1$, both with integer coefficients. $\square$

Figure 1.3. A graph of Berlin.

The deletion–contraction relation (1.1.1) is a natural computing device. For example, the planar graph B in Figure 1.3 that models neighboring districts of Berlin comes with the impressive-looking chromatic polynomial

$$
\begin{aligned}
\chi_B(n) \;=\; & n^{23} - 53\,n^{22} + 1347\,n^{21} - 21845\,n^{20} + 253761\,n^{19} - 2246709\,n^{18} \\
& + 15748804\,n^{17} - 89620273\,n^{16} + 421147417\,n^{15} - 1653474650\,n^{14} \\
& + 5465562591\,n^{13} - 15279141711\,n^{12} + 36185053700\,n^{11} \\
& - 72527020873\,n^{10} + 122562249986\,n^{9} - 173392143021\,n^{8} \\
& + 203081660679\,n^{7} - 193650481777\,n^{6} + 146638574000\,n^{5} \\
& - 84870973704\,n^{4} + 35266136346\,n^{3} - 9362830392\,n^{2} \\
& + 1191566376\,n\,,
\end{aligned} \tag{1.1.2}
$$

which, nevertheless, can be easily computed on any computer. (And yes, $\chi_B(4) = 383904$ is not zero.)

Our proof of Proposition 1.1.1 and, more precisely, the deletion–contraction relation (1.1.1) reveal more about chromatic polynomials, which we invite you to show in Exercise 1.6:

Corollary 1.1.2. *Let G be a loopless graph on $d \geq 1$ nodes and $\chi_G(n) = c_d n^d + c_{d-1} n^{d-1} + \cdots + c_0$ its chromatic polynomial. Then*

(a) *the leading coefficient $c_d = 1$;*
(b) *the constant coefficient $c_0 = 0$;*
(c) $(-1)^d \chi_G(-n) > 0$ *for all integers $n \geq 1$.*

In particular the last property prompts the following natural question which we alluded to in the preface and which lies at the heart of this book.

Do the evaluations $(-1)^{|V|}\chi_G(-n)$ have combinatorial meaning?

This question was first asked (and beautifully answered) by Richard Stanley in 1973. To reproduce his answer, we need the notion of orientations on graphs. Again, to keep the formal pain level at a minimum, we denote the nodes of G by $v_1, v_2, \ldots, v_d$. We define an **orientation** on G through a subset $\rho \subseteq E$; for an edge $e = v_i v_j \in E$ with $i < j$ we direct

$$v_i \overset{e}{\longleftarrow} v_j \text{ if } e \in \rho \qquad \text{and} \qquad v_i \overset{e}{\longrightarrow} v_j \text{ if } e \notin \rho\,.$$

We denote the oriented graph by ${}_\rho G$ and will sometimes write ${}_\rho G = (V, E, \rho)$. Said differently, we may think of G as canonically oriented by directing edges from small index to large, and ρ records the edges on which this orientation is reversed; see Figure 1.4 for an example.

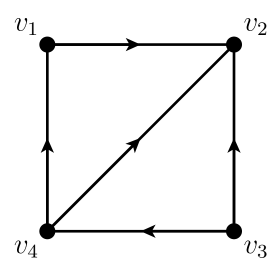

Figure 1.4. An orientation given by $\rho = \{14, 23, 24\}$.

A **directed path** in ${}_\rho G$ is a sequence $v_0, v_1, \ldots, v_s$ of distinct nodes such that $v_{j-1} \to v_j$ is a directed edge in ${}_\rho G$ for all $j = 1, \ldots, s$. If $v_s \to v_0$ is also a directed edge, then $v_0, v_1, \ldots, v_s, v_{s+1} := v_0$ is called a **directed cycle**. An orientation ρ of G is **acyclic** if there are no directed cycles in ${}_\rho G$.

Here is the connection between proper colorings and acyclic orientations: Given a proper coloring c, we define the orientation

$$\rho \ := \ \{v_i v_j \in E \, : \, i < j, \ c(v_i) > c(v_j)\} \, .$$

That is, the edge from lower index i to higher index j is directed along its **color gradient** $c(v_j) - c(v_i)$. We call this orientation ρ **induced** by the coloring c. For example, the orientation pictured in Figure 1.4 is induced by the coloring shown in Figure 1.5.

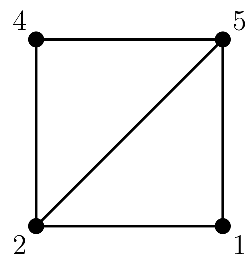

Figure 1.5. A coloring that induces the orientation in Figure 1.4.

Proposition 1.1.3. *Let $c : V \to [n]$ be a proper coloring and ρ the induced orientation on G. Then ${}_\rho G$ is acyclic.*

Proof. Assume that $v_{i_0} \to v_{i_1} \to \cdots \to v_{i_s} \to v_{i_0}$ is a directed cycle in ${}_\rho G$. Then $c(v_{i_0}) < c(v_{i_1}) < \cdots < c(v_{i_s}) < c(v_{i_0})$, which is a contradiction. □

As there are only finitely many acyclic orientations on G, we might count colorings according to the acyclic orientation they induce. An orientation ρ and an n-coloring c of G are called **compatible** if for every oriented edge $u \to v$ in ${}_\rho G$ we have $c(u) \geq c(v)$. The pair (ρ, c) is called **strictly** compatible if $c(u) > c(v)$ for every oriented edge $u \to v$.

Proposition 1.1.4. *If (ρ, c) is strictly compatible, then c is a proper coloring and ρ is an acyclic orientation on G. In particular, $\chi_G(n)$ is the number of strictly compatible pairs (ρ, c), where c is a proper n-coloring.*

Proof. If (ρ, c) are strictly compatible, then, since each edge is oriented, $c(u) > c(v)$ or $c(u) < c(v)$ whenever $uv \in E$. Hence c is a proper coloring and ρ is exactly the orientation induced by c. The acyclicity of ${}_\rho G$ now follows from Proposition 1.1.3. □

We are finally ready for our first combinatorial reciprocity theorem.

Theorem 1.1.5. *Let G be a finite graph on d nodes and $\chi_G(n)$ its chromatic polynomial. Then $(-1)^d\,\chi_G(-n)$ equals the number of compatible pairs (ρ, c), where c is an n-coloring and ρ is an acyclic orientation. In particular, $(-1)^d\,\chi_G(-1)$ equals the number of acyclic orientations of G.*

As one illustration of this theorem, consider the graph G in Figure 1.6; its chromatic polynomial is $\chi_G(n) = n(n-1)(n-2)^2$, and so Theorem 1.1.5 suggests that G should admit 18 acyclic orientations. Indeed, there are six acyclic orientations of the subgraph formed by v_1, v_2, and v_4, and for the remaining two edges, one of the four possible combined orientations of v_2v_3 and v_3v_4 produces a cycle with v_2v_4, so there are a total of $6 \cdot 3 = 18$ acyclic orientations.

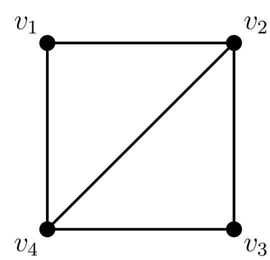

Figure 1.6. This graph has 18 acyclic orientations.

A deletion–contraction proof of Theorem 1.1.5 is outlined in Exercise 1.9; we will give a geometric proof in Section 7.1.

1.2. Flows on Graphs

Given a graph $G = (V, E)$ together with an orientation ρ and the finite Abelian group $\mathbb{Z}_n = \mathbb{Z}/n\mathbb{Z}$, a $\mathbb{Z}_n$**-flow** is a map $f : E \to \mathbb{Z}_n$ that assigns a value $f(e) \in \mathbb{Z}_n$ to each edge $e \in E$ such that there is conservation of flow at every node v:

$$\sum_{\overset{e}{\to} v} f(e) \;=\; \sum_{v \overset{e}{\to}} f(e)\,,$$

that is, what "flows" into the node v is precisely what "flows" out of v. This physical interpretation is a bit shaky as the commodity flowing along edges are elements of $\mathbb{Z}_n$, and the flow conservation is with respect to the group structure. The set

$$\operatorname{supp}(f) \;:=\; \{e \in E \,:\, f(e) \neq 0\}$$

is the **support** of f, and a $\mathbb{Z}_n$-flow f is **nowhere zero** if $\operatorname{supp}(f) = E$. In this section we will be concerned with counting nowhere-zero $\mathbb{Z}_n$-flows, and so we define

$$\varphi_G(n) \ := \ \left|\left\{f \text{ nowhere-zero } \mathbb{Z}_n\text{-flow on } {}_\rho G\right\}\right| .$$

A priori, the counting function $\varphi_G(n)$ depends on our chosen orientation ρ, but our language suggests that this is not the case, which we invite you to verify in Exercise 1.11:

Proposition 1.2.1. *The flow-counting function $\varphi_G(n)$ is independent of the orientation ρ of G.*

A **connected component** of the graph G is a maximal subgraph of G in which any two nodes are connected by a path. A graph G is **connected** if it has only one connected component.[2] As you will discover (at the latest when working on Exercises 1.13 and 1.14), G will not have any nowhere-zero flow if G has a **bridge** (also known as an **isthmus**), that is, an edge whose removal increases the number of connected components of G.

To motivate why we care about counting nowhere-zero flows, we assume that G is a *planar* bridgeless graph with a given embedding into the plane. The drawing of G subdivides the plane into connected regions in which two points lie in the same region whenever they can be joined by a path in $\mathbb{R}^2$ that does not meet G. Two such regions are neighboring if their topological closures share a proper (i.e., 1-dimensional) part of their boundaries. This induces a graph structure on the subdivision of the plane: for the given embedding of G, we define the **dual graph** G^* as the graph with nodes corresponding to the regions and two regions C_1, C_2 share an edge e^* if an original edge e is properly contained in both their boundaries. As we can see in the example pictured in Figure 1.7, the dual graph G^* is typically not simple with parallel edges. If G had bridges, G^* would have loops.

Given an orientation of G, an orientation on G^* is induced by, for example, rotating the edge clockwise. That is, the dual edge will "point" east assuming that the primal edge "points" north:

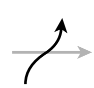

By carefully adding G^* to the picture we can see that dualizing G^* recovers G, i.e., $(G^*)^* = G$.

Our interest in flows lies in the connection to colorings: let c be an n-coloring of G, and for a change we assume that c takes on colors in $\mathbb{Z}_n$. After giving G an orientation, we can record the color gradient $t(uv) = c(v) - c(u)$ for each oriented edge $u \to v$, as shown in Figure 1.8.

[2]These notions refer to an *unoriented* graph.

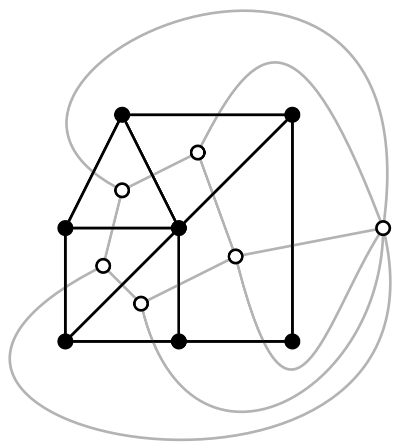

Figure 1.7. A graph and its dual.

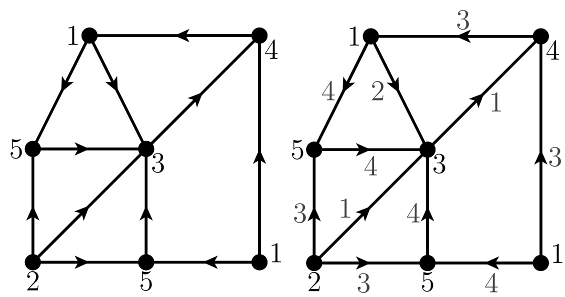

Figure 1.8. Recording color gradients, in $\mathbb{Z}_6$.

Conversely, knowing the color of a single node v_0, we can recover the coloring from $t : E \to \mathbb{Z}_n$: for a node $v \in V$ simply choose an undirected path $v_0 = p_0 p_1 p_2 \cdots p_k = v$ from v_0 to v. Then while walking along this path we can color each node p_i by adding or subtracting $t(p_{i-1}p_i)$ to $c(p_{i-1})$ depending on whether we walked the edge $p_{i-1}p_i$ with or against its orientation.

The color $c(v)$ is independent of the chosen path and thus, walking along a cycle in G the sum of the values $t(e)$ of edges along their orientation minus those against their orientation has to be zero; this is illustrated in Figure 1.9. Now, via the correspondence of primal and dual edges, t induces a map $f : E^* \to \mathbb{Z}_n$ on the dual graph G^*, shown in Figure 1.10. Each node of G^* represents a region that is bounded by a cycle in G, and the orientation on G^* is such that walking around this cycle clockwise, each edge traversed along its orientation corresponds to a dual edge into the

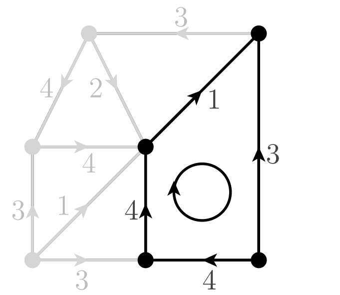

$$4 + 1 - 3 + 4 \equiv 0 \pmod 6$$

Figure 1.9. A cycle of flows.

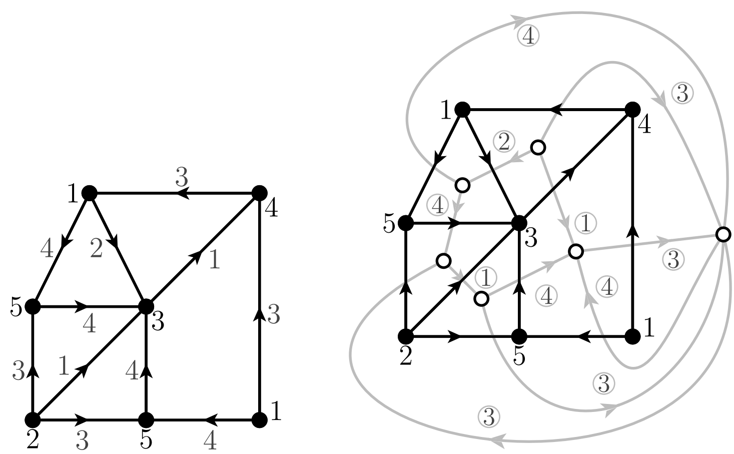

Figure 1.10. A flow and its dual.

region while counter-clockwise edges dually point out of the region. The cycle condition, illustrated in Figure 1.11, then proves:

Proposition 1.2.2. *Let G be a connected planar graph with dual G^*. For every n-coloring c of G, the induced map f is a $\mathbb{Z}_n$-flow on G^*, and every such flow arises this way. Moreover, the coloring c is proper if and only if f is nowhere zero.*

Conversely, for a given (nowhere-zero) flow f on G^* one can construct a (proper) coloring on G (see Exercise 1.12). In light of all this, we can rephrase the Four-color Theorem as follows.

Corollary 1.2.3 (Dual Four-color Theorem)**.** *If G is a planar bridgeless graph, then $\varphi_G(4) > 0$.*

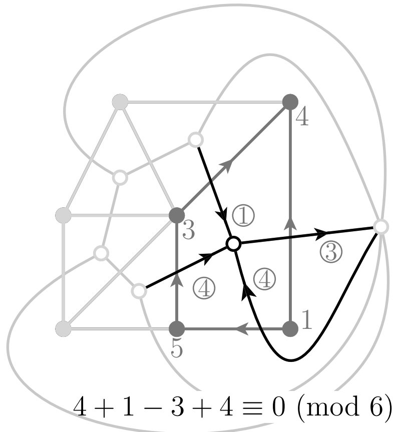

Figure 1.11. Proposition 1.2.2 illustrated.

This perspective on colorings of planar graphs was pioneered by William Tutte who initiated the study of $\varphi_G(n)$ for all (not necessarily planar) graphs. To see how much flows differ from colorings, we observe that there is no universal constant n_0 such that every graph has a proper n_0-coloring. The analogous statement for flows is not so clear and, in fact, Tutte conjectured the following:

Five-flow Conjecture. *Every bridgeless graph has a nowhere-zero $\mathbb{Z}_5$-flow.*

This sounds like a rather daring conjecture, as it is not even clear that there is *any* n such that every bridgeless graph has a nowhere-zero $\mathbb{Z}_n$-flow. However, it was shown by Paul Seymour that $n \le 6$ works. In Exercise 1.17 you will show that there exist graphs that do not admit a nowhere-zero $\mathbb{Z}_4$-flow.

On the enumerative side, we have the following.

Proposition 1.2.4. *If G is a bridgeless connected graph, then $\varphi_G(n)$ agrees with a polynomial with integer coefficients of degree $|E|-|V|+1$ and leading coefficient* 1*.*

Again, we will abuse notation and refer to $\varphi_G(n)$ as the **flow polynomial** of G. The proof of the polynomiality is a deletion–contraction argument which is deferred to Exercise 1.13.

Towards a reciprocity statement, we need a notion dual to acyclic orientations: an orientation ρ on G is **totally cyclic** if every edge in ${}_\rho G$ is contained in a directed cycle. We quickly define the **cyclotomic number** of G as $\xi(G) := |E| - |V| + c$, where $c = c(G)$ is the number of connected components of G.

Theorem 1.2.5. *Let G be a bridgeless graph. For every positive integer n, the evaluation $(-1)^{\xi(G)}\varphi_G(-n)$ counts the number of pairs (f, ρ), where f is a $\mathbb{Z}_n$-flow and ρ is a totally-cyclic reorientation of $G/\operatorname{supp}(f)$. In particular, $(-1)^{\xi(G)}\varphi_G(-1)$ equals the number of totally-cyclic orientations of G.*

We will prove this theorem in Section 7.6.

1.3. Order Polynomials

A **partially ordered set**, or **poset** for short, is a set Π together with a binary relation $\preceq_\Pi$ that is

reflexive:	$a \preceq_\Pi a$,
transitive:	$a \preceq_\Pi b \preceq_\Pi c$ implies $a \preceq_\Pi c$, and
antisymmetric:	$a \preceq_\Pi b$ and $b \preceq_\Pi a$ implies $a = b$

for all $a, b, c \in \Pi$. We write $\preceq$ if the poset is clear from the context.

Partially ordered sets are ubiquitous structures in combinatorics and, as we will amply demonstrate soon, are indispensable in enumerative and geometric combinatorics. Most posets that we will encounter in this book are finite and when we say *poset*, we will always mean a finite poset unless stated otherwise.

The essence of a poset is encoded by its **cover relations**: an element $a \in \Pi$ is covered by an element b if

$$[a, b] \;:=\; \{z \in \Pi \,:\, a \preceq z \preceq b\} \;=\; \{a, b\}\,;$$

in plain English: $a \prec b$ and there is nothing between a and b. We write $a \prec\!\cdot\, b$ when a is covered by b. From its cover relations we can recover the poset by taking the transitive closure and adding in the reflexive relations. The cover relations can be thought of as a directed graph, and this gives an effective way to picture a poset: The **Hasse diagram** of Π is a drawing of the directed graph of cover relations in Π as an (undirected) graph where the node a is drawn lower than the node b whenever $a \prec b$. Here is an example: for $n \in \mathbb{Z}_{>0}$ we define D_n as the set $[n] = \{1, 2, \dots, n\}$ ordered by divisibility, that is, $a \preceq b$ if a divides b. The Hasse diagram of D_{10} is given in Figure 1.12.

This example truly is a *partial* order as, for example, 2 and 7 are not comparable. A poset in which each element is comparable to every other element is a **chain**. To be more precise: the poset Π is a chain if we have either $a \preceq b$ or $b \preceq a$ for any two elements $a, b \in \Pi$. The elements of a chain are **totally** or **linearly ordered**.

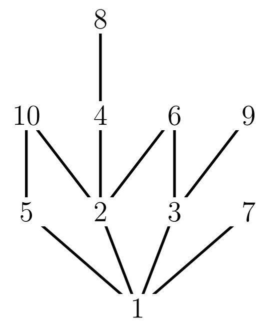

Figure 1.12. D_{10}: the set [10], partially ordered by divisibility.

A map $\phi : \Pi \to \Pi'$ is **(weakly) order preserving** if for all $a, b \in \Pi$

$$a \preceq_\Pi b \quad \Longrightarrow \quad \phi(a) \preceq_{\Pi'} \phi(b)$$

and **strictly order preserving** if

$$a \prec_\Pi b \quad \Longrightarrow \quad \phi(a) \prec_{\Pi'} \phi(b) \,.$$

For example, we can label the elements of a chain Π such that

$$\Pi \; = \; \{a_1 \prec a_2 \prec \cdots \prec a_n\} \,,$$

which makes Π **isomorphic** to $[n] = \{1 < 2 < \cdots < n\}$, in the sense that there is a bijection $\phi : \Pi \to [n]$ such that ϕ and ϕ^{-1} are strictly order preserving.

Order-preserving maps are the natural morphisms (even in a categorical sense) between posets, and in this section we will be concerned with counting (strictly) order-preserving maps from a poset into chains.

A strictly order-preserving map ϕ from one chain $[d]$ into another $[n]$ exists only if $d \le n$ and is then determined by

$$1 \; \le \; \phi(1) \; < \; \phi(2) \; < \; \cdots \; < \; \phi(d) \; \le \; n \,.$$

Thus, the number of such maps equals $\binom{n}{d}$, the number of d-subsets of an n-set. In the case of a general poset Π, we define the **strict order polynomial**

$$\Omega^\circ_\Pi(n) \; := \; |\{\phi : \Pi \to [n] \text{ strictly order preserving}\}| \,.$$

As we have just seen, $\Omega^\circ_\Pi(n)$ is indeed a polynomial when $\Pi = [d]$. We now show that polynomiality holds for all posets Π.

Proposition 1.3.1. *For a finite poset* Π*, the function* $\Omega^\circ_\Pi(n)$ *agrees with a polynomial of degree* $|\Pi|$ *with rational coefficients.*

Proof. Let $d := |\Pi|$ and $\phi : \Pi \to [n]$ be a strictly order-preserving map. Now ϕ factors uniquely into a surjective map σ onto $\phi(\Pi)$ followed by an injection ι:

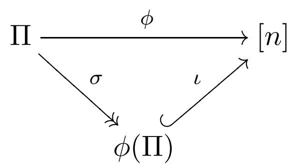

(Use the functions $\sigma(a) := \phi(a)$ and $\iota(a) := a$, defined with domains and codomains pictured above.) The image $\phi(\Pi)$ is a subposet of a chain and so is itself a chain. Thus $\Omega^\circ_\Pi(n)$ counts the number of pairs (σ, ι) of strictly order-preserving maps $\Pi \twoheadrightarrow [r] \rightarrowtail [n]$ for $r = 1, 2, \dots, d$. For fixed r, there are only finitely many order-preserving surjections $\sigma : \Pi \to [r]$, say, s_r many. As we discussed earlier, the number of strictly order-preserving maps $[r] \to [n]$ is exactly $\binom{n}{r}$, which is a rational polynomial in n of degree r. Hence, for fixed r, there are $s_r \binom{n}{r}$ many pairs (σ, ι) and we obtain

$$\Omega^\circ_\Pi(n) \;=\; s_d \binom{n}{d} + s_{d-1} \binom{n}{d-1} + \cdots + s_1 \binom{n}{1},$$

which finishes our proof. □

As an aside, Proposition 1.3.1 proves that $\Omega^\circ_\Pi(n)$ is a polynomial with *integral* coefficients if we use $\left\{ \binom{n}{r} \;:\; r \in \mathbb{Z}_{\geq 0} \right\}$ as a basis for the polynomial ring $\mathbb{R}[n]$. That the binomial coefficients indeed form a basis for the univariate polynomials follows from Proposition 1.3.1: if Π is an **antichain** on d elements, i.e., a poset in which no elements are related, then

$$\Omega^\circ_\Pi(n) \;=\; n^d \;=\; s_d \binom{n}{d} + s_{d-1} \binom{n}{d-1} + \cdots + s_1 \binom{n}{1}. \tag{1.3.1}$$

In this case, the coefficients $s_r = S(d, r)$ are the **Stirling numbers of the second kind** which count the number of surjective maps $[d] \twoheadrightarrow [r]$. (The Stirling numbers might come in handy in Exercise 1.10.)

For the case that Π is a d-chain, the reciprocity statement (0.0.2) says that $(-1)^d\, \Omega^\circ_\Pi(-n)$ gives the number of d-multisubsets of an n-set, which equals, in turn, the number of (weak) order-preserving maps from a d-chain to an n-chain. Our next combinatorial reciprocity theorem expresses this duality between weak and strict order-preserving maps from a general poset into chains. You can already guess what is coming. We define the **order polynomial**

$$\Omega_\Pi(n) \;:=\; |\{\phi : \Pi \to [n] \text{ order preserving}\}|.$$

A slight modification (which we invite you to check in Exercise 1.20) of our proof of Proposition 1.3.1 implies that $\Omega_\Pi(n)$ indeed agrees with a polynomial in n of degree $|\Pi|$, and the following reciprocity theorem gives the relationship between the two polynomials $\Omega_\Pi(n)$ and $\Omega_\Pi^\circ(n)$.

Theorem 1.3.2. *Let Π be a finite poset. Then*

$$(-1)^{|\Pi|}\,\Omega_\Pi^\circ(-n) \;=\; \Omega_\Pi(n)\,.$$

We will prove this theorem in Chapter 2. To further motivate the study of order polynomials, we remark that a poset Π gives rise to an oriented graph by way of the cover relations of Π. Conversely, the binary relation given by an oriented graph G can be completed to a partial order $\Pi(G)$ by adding the necessary transitive and reflexive relations if and only if G is acyclic. Figure 1.13 shows an example, for the orientation pictured in Figure 1.4. The following result will be the subject of Exercise 1.18.

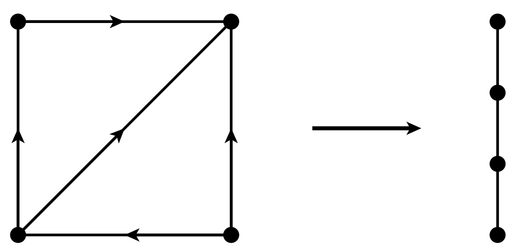

Figure 1.13. From an acyclic orientation to a poset.

Proposition 1.3.3. *Let ${}_\rho G = (V, E, \rho)$ be an acyclic graph and $\Pi({}_\rho G)$ the induced poset. A map $c : V \to [n]$ is strictly compatible with the orientation ρ of G if and only if c is a strictly order-preserving map $\Pi({}_\rho G) \to [n]$.*

In Proposition 1.1.4 we identified the number of n-colorings $\chi_G(n)$ of G as the number of colorings c strictly compatible with some acyclic orientation ρ of G, and so this proves:

Corollary 1.3.4. *The chromatic polynomial $\chi_G(n)$ of a graph G is the sum of the order polynomials $\Omega^\circ_{\Pi({}_\rho G)}(n)$ for all acyclic orientations ρ of G.*

1.4. Ehrhart Polynomials

The formulation of (0.0.1) in terms of d-subsets of an n-set has a straightforward geometric interpretation that will fuel much of what is about to come: the d-subsets of $[n]$ correspond precisely to the points in $\mathbb{R}^d$ with integral coordinates in the set

$$(n+1)\triangle_d^\circ \;=\; \left\{\mathbf{x} \in \mathbb{R}^d \,:\, 0 < x_1 < x_2 < \cdots < x_d < n+1\right\}. \tag{1.4.1}$$

Next we explain the notation on the left-hand side: we define

$$\triangle_d^\circ := \left\{ \mathbf{x} \in \mathbb{R}^d : 0 < x_1 < x_2 < \cdots < x_d < 1 \right\},$$

and for a set $S \subseteq \mathbb{R}^d$ and a positive integer n, we set

$$n\,S := \{ n\,\mathbf{x} : \mathbf{x} \in S \},$$

the **n-th dilate** of S. (We hope the notation in (1.4.1) now makes sense.) For example, when $d = 2$,

$$\triangle_2^\circ = \left\{ (x_1, x_2) \in \mathbb{R}^2 : 0 < x_1 < x_2 < 1 \right\}$$

is the interior of a triangle, and every integer point (x_1, x_2) in the $(n+1)$-st dilate of $\triangle_2^\circ$ satisfies $0 < x_1 < x_2 < n+1$ or, equivalently, $1 \le x_1 < x_2 \le n$. We illustrate these integer points for the case $n = 5$ in Figure 1.14.

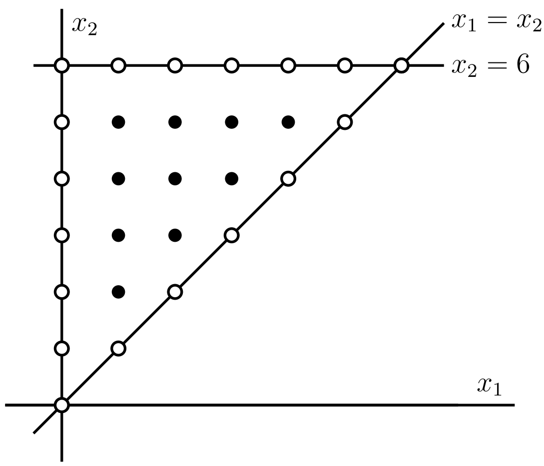

Figure 1.14. The integer points in $6\triangle_2^\circ$.

A **convex lattice polygon** $\mathsf{P} \subset \mathbb{R}^2$ is the smallest convex set containing a given finite set of noncollinear integer points in the plane. The **interior** of P is denoted by P°. Convex polygons are 2-dimensional instances of **convex polytopes**, which live in any dimension and whose properties we will study in detail in Chapter 3. For now, we count on your intuition about terms like *convex* and objects such as *vertices* and *edges* of a polygon, which will be defined rigorously in Chapter 3.

For a bounded set $S \subset \mathbb{R}^2$, we write $E(S) := \left| S \cap \mathbb{Z}^2 \right|$ for the number of integer lattice points in S. Our example above motivates the definitions of the counting functions

$$\operatorname{ehr}_{\mathsf{P}^\circ}(n) := E(n\,\mathsf{P}^\circ) = \left| n\,\mathsf{P}^\circ \cap \mathbb{Z}^2 \right|$$

and

$$\operatorname{ehr}_{\mathsf{P}}(n) \;:=\; E(n\,\mathsf{P}) \;=\; \left|n\,\mathsf{P} \cap \mathbb{Z}^2\right|,$$

the **Ehrhart functions** of P. The historical reasons for this naming convention will be given in Chapters 4 and 5.

As we know from (0.0.1), the number of integer lattice points in the $(n+1)$-st dilate of $\triangle_2^\circ$ is given by the polynomial

$$\operatorname{ehr}_{\triangle_2^\circ}(n+1) \;=\; \binom{n}{2}.$$

To make the combinatorial reciprocity statement given by (0.0.1) geometric, we observe that the number of weak order-preserving maps from $[n]$ into $[2]$ is given by the integer points in the $(n-1)$-st dilate of

$$\triangle_2 \;=\; \left\{(x_1, x_2) \in \mathbb{R}^2 \,:\, 0 \le x_1 \le x_2 \le 1\right\},$$

the closure of $\triangle_2^\circ$. The combinatorial reciprocity statement given by (0.0.1) now reads $(-1)^2\binom{-n}{2}$ *equals the number of integer points in* $(n-1)\triangle_2$. Unraveling the parameters (and making appropriate shifts), we can rephrase this as: $(-1)^2 \operatorname{ehr}_{\triangle_2^\circ}(-n)$ equals the number of integer points in $n\triangle_2$. The reciprocity theorem featured in this section states that this holds for all convex lattice polygons; in Chapter 5 we will prove an analogue in all dimensions.

Theorem 1.4.1. *Let* $\mathsf{P} \subset \mathbb{R}^2$ *be a lattice polygon. Then* $\operatorname{ehr}_{\mathsf{P}}(n)$ *agrees with a polynomial of degree* 2 *with rational coefficients, and* $(-1)^2 \operatorname{ehr}_{\mathsf{P}}(-n)$ *equals the number of integer points in* $n\mathsf{P}^\circ$.

In the remainder of this section we will prove this theorem. The proof will be a series of simplifying steps that are similar in spirit to those that we will employ for the general result in Section 5.2.

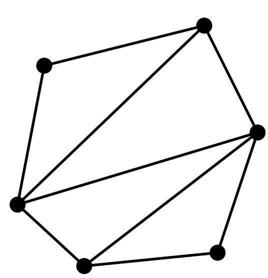

Figure 1.15. A triangulation of a hexagon.

As a first step, we reduce the problem of showing polynomiality for Ehrhart functions of arbitrary lattice polygons to that of lattice *triangles*.

Let P be a lattice polygon in the plane with n vertices. We can **triangulate** P by cutting the polygon along sufficiently many (exactly $n-3$) nonintersecting diagonals, as in Figure 1.15. The result is a set of $n-2$ lattice triangles that cover P. We denote by $\mathcal{T}$ the collection of **faces** of all these triangles, that is, $\mathcal{T}$ consists of n zero-dimensional polytopes (vertices), $2n-3$ one-dimensional polytopes (edges), and $n-2$ two-dimensional polytopes (triangles).

Our triangulation is a well-behaved collection of polytopes in the plane in the sense that they intersect nicely: if two elements of $\mathcal{T}$ intersect, then they intersect in a common face of both. This is useful, as counting lattice points is a *valuation*.[3] Namely, for $S, T \subset \mathbb{R}^2$,

$$E(S \cup T) \;=\; E(S) + E(T) - E(S \cap T)\,, \tag{1.4.2}$$

and applying the inclusion–exclusion relation (1.4.2) repeatedly to the elements in our triangulation of P yields

$$\operatorname{ehr}_{\mathsf{P}}(n) \;=\; \sum_{\mathsf{F}\in\mathcal{T}} \mu(\mathsf{F})\operatorname{ehr}_{\mathsf{F}}(n)\,, \tag{1.4.3}$$

where the $\mu(\mathsf{F})$ are some coefficients that correct for over-counting. If F is a triangle, then $\mu(\mathsf{F}) = 1$—after all, we want to count the lattice points in P that are covered by the triangles. For an edge F of the triangulation, we have to make the following distinction: F is an **interior edge** of $\mathcal{T}$ if it is contained in two triangles. In this case the lattice points in F get counted twice, and in order to compensate for this, we set $\mu(\mathsf{F}) = -1$. In the case that F is a **boundary edge**, i.e., F lies in only one triangle of $\mathcal{T}$, there is no over-counting and we can set $\mu(\mathsf{F}) = 0$. To generalize this to all faces of $\mathcal{T}$, we call a face $\mathsf{F} \in \mathcal{T}$ a **boundary face** of $\mathcal{T}$ if F is contained in the boundary of P, and an **interior face** otherwise. We can give the coefficients $\mu(\mathsf{F})$ explicitly as follows.

Proposition 1.4.2. *Let $\mathcal{T}$ be a triangulation of a lattice polygon $\mathsf{P} \subset \mathbb{R}^2$. Then the coefficients $\mu(\mathsf{F})$ in (1.4.3) are given by*

$$\mu(\mathsf{F}) \;=\; \begin{cases} (-1)^{2-\dim \mathsf{F}} & \text{if } \mathsf{F} \text{ is interior,} \\ 0 & \text{otherwise.} \end{cases}$$

For boundary vertices $\mathsf{F} = \{\mathbf{v}\}$, we can check that $\mu(\mathsf{F}) = 0$ is correct: the vertex is counted positively as a lattice point by every incident triangle and negatively by every incident interior edge. As there are exactly one interior edge less than incident triangles, we do not count the vertex more than once. For an interior vertex, the number of incident triangles and incident (interior) edges are equal and hence $\mu(\mathsf{F}) = 1$. (In triangulations

[3] We'll have more to say about valuations in Section 3.4.

of P obtained by cutting along diagonals we never encounter *interior* vertices, however, they will appear soon when we consider a different type of triangulation.)

The coefficient $\mu(\mathsf{F})$ for a triangulation of a polygon was easy to argue and to verify in the plane. For higher-dimensional polytopes we will have to resort to more algebraic and geometric means. The right algebraic setup will be discussed in Chapter 2 where we will make use of the fact that a triangulation $\mathcal{T}$ constitutes a partially ordered set. In the language of posets, $\mu(\mathsf{F})$ is an evaluation of the *Möbius function* for the poset $\mathcal{T}$. Möbius functions are esthetically satisfying but are in general difficult to compute. However, we are dealing with situations with plenty of geometry involved, and we will make use of that in Chapter 5 to give a statement analogous to Proposition 1.4.2 in general dimension.

Returning to our 2-dimensional setting, showing that $\mathrm{ehr}_{\mathsf{F}}(n)$ is a polynomial whenever F is a lattice point, a lattice segment, or a lattice triangle gives us the first half of Theorem 1.4.1. If F is a vertex, then $\mathrm{ehr}_{\mathsf{F}}(n) = 1$. If $\mathsf{F} \in \mathcal{T}$ is an edge of one of the triangles and thus a lattice segment, verifying that $\mathrm{ehr}_{\mathsf{F}}(n)$ is a polynomial is the content of Exercise 1.21.

The remaining challenge now is the polynomiality and reciprocity for lattice triangles. For the rest of this section, let $\triangle \subset \mathbb{R}^2$ be a fixed lattice triangle in the plane. The idea that we will use is to triangulate the dilates $n\triangle$ for $n \geq 1$, but the triangulation will change with n. Figure 1.16 gives the picture for $n = 1, 2, 3$.

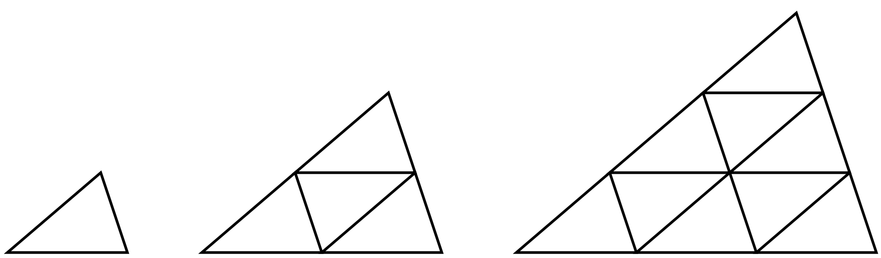

Figure 1.16. Special triangulations of dilates of a lattice triangle.

We trust that you can imagine the triangulation for all values of n. The special property of this triangulation is that up to lattice translations, there are only a few different pieces. In fact, there are only two different lattice triangles used in the triangulation of $n\triangle$: there is $\triangle$ itself and (lattice translates of) the reflection of $\triangle$ with respect to the origin, which we will denote by $\triangledown$. As for edges, we have three different kinds of edges, namely, the edges $\diagdown$, —, and $\diagup$. Up to lattice translation, there is only one vertex $\bullet$.

Now we count how many copies of each tile occur in these special triangulations; let $t(\mathsf{Q}, n)$ denote the number of times Q appears in our triangulation of $n\triangle$. For triangles, we count

$$t(\triangle, n) \;=\; \binom{n+1}{2} \quad \text{and} \quad t(\triangledown, n) \;=\; \binom{n}{2}.$$

For the interior edges, we observe that each interior edge is incident to a unique upside-down triangle $\triangledown$ and consequently

$$t(\diagdown, n) \;=\; t(—, n) \;=\; t(\diagup, n) \;=\; \binom{n}{2}.$$

Similarly, for interior vertices,

$$t(\bullet, n) \;=\; \binom{n-1}{2}.$$

Thus with (1.4.3), the Ehrhart function for the triangle $\triangle$ is

$$\begin{aligned} \operatorname{ehr}_{\triangle}(n) \;=\; & \binom{n+1}{2} E(\triangle) + \binom{n}{2} E(\triangledown) \\ & - \binom{n}{2} \big(E(\diagdown) + E(—) + E(\diagup)\big) \\ & + \binom{n-1}{2} E(\bullet)\,. \end{aligned} \tag{1.4.4}$$

This proves that $\operatorname{ehr}_{\triangle}(n)$ agrees with a polynomial of degree 2, and together with (1.4.3) this establishes the first half of Theorem 1.4.1.

To prove the combinatorial reciprocity of Ehrhart polynomials in the plane, we make the following useful observation.

Proposition 1.4.3. *If for every lattice polygon* $\mathsf{P} \subset \mathbb{R}^2$ *we have that* $\operatorname{ehr}_{\mathsf{P}}(-1)$ *equals the number of lattice points in the interior of* P*, then* $\operatorname{ehr}_{\mathsf{P}}(-n) = E(n\mathsf{P}^\circ)$ *for all* $n \ge 1$.

Proof. For fixed $n \ge 1$, we denote by Q the lattice polygon $n\mathsf{P}$. We see that $\operatorname{ehr}_{\mathsf{Q}}(m) = E(m(n\mathsf{P}))$ for all $m \ge 1$. Hence the Ehrhart polynomial of Q is given by $\operatorname{ehr}_{\mathsf{P}}(mn)$ and for $m = -1$ we conclude

$$\operatorname{ehr}_{\mathsf{P}}(-n) \;=\; \operatorname{ehr}_{\mathsf{Q}}(-1) \;=\; E(\mathsf{Q}^\circ) \;=\; E(n\mathsf{P}^\circ)\,,$$

which finishes our proof. □

To establish the combinatorial reciprocity of Theorem 1.4.1 for triangles, we can simply substitute $n = -1$ into (1.4.4) and use (0.0.2) to obtain

$$\operatorname{ehr}_{\triangle}(-1) \;=\; E(\triangledown) - E(\diagdown) - E(—) - E(\diagup) + 3E(\bullet)\,,$$

which equals the number of interior lattice points of $\triangledown$. Observing that $\triangle$ and $\triangledown$ have the same number of lattice points finishes the argument.

For the general case, Exercise 1.21 gives

$$\mathrm{ehr}_{\mathsf{P}}(-1) \;=\; \sum_{\mathsf{F}\in\mathcal{T}} E(\mathsf{F}^\circ) \;=\; E(\mathsf{P}^\circ)$$

and this (finally!) concludes our proof of Theorem 1.4.1.

Exercises 1.21 and 1.23 also answer the question of why we carefully triangulate P along diagonals (as opposed to cutting it up arbitrarily to obtain triangles): Theorem 1.4.1 is only true for lattice polygons. There are versions for polygons with rational and irrational coordinates but they become increasingly complicated. By cutting along diagonals we can decompose a lattice polygon into lattice segments and lattice triangles. This part becomes nontrivial already in dimension 3, and we will worry about this in Chapter 5.

In Exercise 1.25 we will look into the question as to what the coefficients of $\mathrm{ehr}_{\mathsf{P}}(n)$, for a lattice polygon P, tell us. We finish this chapter by considering the constant coefficient $c_0 = \mathrm{ehr}_{\mathsf{P}}(0)$. This is the most tricky one, as we could argue that $\mathrm{ehr}_{\mathsf{P}}(0) = E(0\mathsf{P})$ and since $0\mathsf{P}$ is just a single point, we get $c_0 = 1$. This argument is flawed: we defined $\mathrm{ehr}_{\mathsf{P}}(n)$ only for $n \geq 1$. To see that this argument is, in fact, plainly wrong, we consider $S = \mathsf{P}_1 \cup \mathsf{P}_2 \subset \mathbb{R}^2$, where P_1 and P_2 are disjoint lattice polygons. Since they are disjoint, $\mathrm{ehr}_S(n) = \mathrm{ehr}_{\mathsf{P}_1}(n) + \mathrm{ehr}_{\mathsf{P}_2}(n)$. Now $0S$ is also just a point and therefore

$$1 \;=\; \mathrm{ehr}_S(0) \;=\; \mathrm{ehr}_{\mathsf{P}_1}(0) + \mathrm{ehr}_{\mathsf{P}_2}(0) \;=\; 2\,.$$

It turns out that $c_0 = 1$ is still correct but the justification will have to wait until Theorem 5.1.8. In Exercise 1.26, you will prove a more general version for Theorem 1.4.1 that dispenses of convexity.

Notes

Graph-coloring problems started in the form of coloring maps such that countries sharing a proper part of their boundaries get colored with different colors. The graphs associated to such map-coloring problems are planar as is illustrated in Figure 1.3. So the fact that the chromatic polynomial is indeed a polynomial was proved for maps (in 1912 by George Birkhoff [**32**]) before Hassler Whitney proved it for general graphs in 1932 [**184**]. The deletion–contraction argument that we used in the proof of Proposition 1.1.1 gives an algorithm that we used, for example, for the chromatic polynomial (1.1.2) of Berlin. Complexity-theory-savvy readers might want to ponder the (exponential) complexity of this algorithm but it can be implemented with little effort (we used SAGE [**55**]) and for small graphs it works well. As we mentioned, the first proof of the Four-color Theorem is due to Kenneth Appel and Wolfgang Haken [**7**, **8**]. Theorem 1.1.5 is due to

Richard Stanley [**161**]. We will give a proof from a geometric point of view in Section 7.1.

As already mentioned, the approach of studying colorings of planar graphs through flows on their duals was pioneered by William Tutte [**179**], who also conceived the Five-flow Conjecture. This conjecture becomes a theorem when “5” is replaced by “6”, due to Paul Seymour [**154**]; the 8-flow theorem had previously been shown by François Jaeger [**93, 94**]. Theorem 1.2.5 was proved in [**37**]. We will give a proof in Section 7.6.

The number of proper n-colorings, of nowhere-zero $\mathbb{Z}_n$-flows, and of acyclic or totally cyclic orientations can all be computed by using deletions and contractions. More generally, let f be a function that assigns any graph G a number $f(G) \in \mathbb{R}$ such that $f(G) = f(G')$ if G and G' are isomorphic. Then f is called a *generalized Tutte–Grothendieck invariant* if there are constants α, β such that for any $e \in E(G)$

$$f(G) \;=\; \begin{cases} \alpha\, f(G \setminus e) + \beta\, f(G/e) & \text{if } e \text{ is neither a loop nor a bridge,} \\ f(e)\, f(G \setminus e) & \text{otherwise.} \end{cases}$$

Here $f(e)$ is the value on the graph that consists of the edge e alone. It is not difficult to show that there is a *universal* Tutte–Grothendieck invariant in the following sense: for every graph G there is a polynomial $T_G(x, y) \in \mathbb{Z}[x, y]$ such that $f(G)$ is an evaluation of $T_G(x, y)$ in terms of α, β, and the values of f on a loop and bridge; see [**44**] for much more on this. The polynomial $T_G(x, y)$ is called the *Tutte polynomial* of G. Its evaluations, its coefficients, as well as the many mathematical contexts in which they occur are quite remarkable, and that area of geometric and algebraic combinatorics is very active. We will see the notion of deletion–contraction in a more geometric context in Chapter 7.

Order polynomials were introduced by Richard Stanley [**160, 166**] as “chromatic-like polynomials for posets” (this is reflected in Corollary 1.3.4); Theorem 1.3.2 is due to him. We will study order polynomials in depth in Chapters 2 and 6.

Theorem 1.4.1 is essentially due to Georg Pick [**136**], whose famous formula is the subject of Exercise 1.25. In some sense, this formula marks the beginning of the study of integer-point enumeration in polytopes. Our phrasing of Theorem 1.4.1 suggests that it has an analogue in higher dimensions, and we will study this analogue in Chapters 4 and 5.

Herbert Wilf [**185**] raised the question of characterizing which polynomials can occur as chromatic polynomials of graphs. This question has spawned a lot of work in algebraic combinatorics. For example, a recent theorem of June Huh [**89**] says that the absolute values of the coefficients of every chromatic polynomial form a *unimodal* sequence, that is, the sequence increases up to some point, after which it decreases. Huh's theorem had been conjectured by Ronald Read [**140**] almost 50 years earlier. In fact, Huh proved much more. In Chapter 7 we will study arrangements of hyperplanes and their associated characteristic polynomials. Huh and later Huh and Eric Katz [**90**] proved that, up to sign, the coefficients of characteristic polynomials of hyperplane arrangements (defined over any field) form a *log-concave* sequence. We will see the relation between chromatic and characteristic polynomials in Chapter 7.

Exercises

1.1 Two graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ are **isomorphic** if there is a bijection $\phi : V_1 \to V_2$ such that for all $u, v \in V_1$

$$uv \in E_1 \quad \text{if and only if} \quad \phi(u)\,\phi(v) \in E_2.$$

Let G be a planar graph and let G_1 and G_2 be the dual graphs for two distinct planar embeddings of G. Is it true that G_1 and G_2 are isomorphic?

If not, can you give a sufficient condition on G such that the above claim is true? (*Hint:* A precise characterization is rather difficult, but for a sufficient condition you might want to contemplate Steinitz's theorem [**176**]; see [**190**, Ch. 4] for a modern treatment.)

1.2 Find two simple nonisomorphic graphs G and H with $\chi_G(n) = \chi_H(n)$. Can you find many (polynomial, exponential) such examples in the number of nodes? Can you make your examples arbitrarily high connected?

1.3 Find the chromatic polynomials of
   (a) the path on d nodes;
   (b) the cycle on d nodes;
   (c) the wheel with d spokes (and $d+1$ nodes); for example, the wheel with six spokes is this:

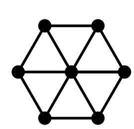

1.4 Verify that the graph of Berlin in Figure 1.3 cannot be colored with three colors. (*Hint:* Instead of evaluating the chromatic polynomial, try to find a simple subgraph that is not 3-colorable.)

1.5 Show that if G has c connected components, then n^c divides the polynomial $\chi_G(n)$.

1.6 ⬠ Complete the proof of Corollary 1.1.2: Let G be a loopless nonempty graph on d nodes and $\chi_G(n) = c_d n^d + c_{d-1} n^{d-1} + \cdots + c_0$ its chromatic polynomial. Then
   (a) the leading coefficient $c_d = 1$;
   (b) the constant coefficient $c_0 = 0$;
   (c) $(-1)^d \chi_G(-n) > 0$.

1.7 Prove that every **complete graph** K_d (a graph with d nodes and all possible edges between them) has exactly $d!$ acyclic orientations.

1.8 Using a construction similar to the one in our proof of Proposition 1.3.1, show that the chromatic polynomial of a given graph G can be written as

$$\chi_G(n) \;=\; a_d \binom{n}{d} + a_{d-1} \binom{n}{d-1} + \cdots + a_1 \binom{n}{1}$$

for some (explicitly describable) nonnegative integers $a_1, a_2, \ldots, a_d$. (This gives yet another proof of Proposition 1.1.1.)

1.9 In this exercise you will give a deletion–contraction proof of Theorem 1.1.5.
   (a) Verify that the deletion–contraction relation (1.1.1) implies for the function $\overline{\chi}_G(n) := (-1)^d \chi_G(-n)$ that

$$\overline{\chi}_G(n) \;=\; \overline{\chi}_{G \setminus e}(n) + \overline{\chi}_{G/e}(n) \,.$$

   (b) Define $\mathcal{X}_G(n)$ as the number of compatible pairs of an acyclic orientation ρ and an n-coloring c. Show $\mathcal{X}_G(n)$ satisfies the same deletion–contraction relation as $\overline{\chi}_G(n)$.
   (c) Infer that $\overline{\chi}_G(n) = \mathcal{X}_G(n)$ by induction on $|E|$.

1.10 The **complete bipartite graph** $K_{r,s}$ is the graph on the node set $V = \{1, 2, \ldots, r, 1', 2', \ldots, s'\}$ and edges

$$E \;=\; \left\{ ij' \,:\, 1 \le i \le r,\ 1 \le j \le s \right\}.$$

Determine the chromatic polynomial $\chi_{K_{r,s}}(n)$ for $m, k \ge 1$. (*Hint:* Proper n-colorings of $K_{r,s}$ correspond to pairs (f, g) of maps $f : [r] \to [n]$ and $g : [s] \to [n]$ with disjoint ranges.)

1.11 ⬠ Prove Proposition 1.2.1: The flow-counting function $\varphi_G(n)$ is independent on the orientation of G.

1.12 ⬠ Let G be a connected planar graph with dual G^*. By reversing the steps in our proof before Proposition 1.2.2, show that every (nowhere-zero) $\mathbb{Z}_n$-flow f on G^* naturally gives rise to n different (proper) n-colorings on G.

1.13 ⬠ Prove Proposition 1.2.4: If G is a bridgeless connected graph, then $\varphi_G(n)$ agrees with a monic polynomial of degree $|E| - |V| + 1$ with integer coefficients.

1.14 ⬠ Let $G = (V, E)$ be a graph, and let n be a positive integer. An **n-flow** is a function $g : E \to \mathbb{Z}$ with $-n < g(e) < n$ such that conservation of flow holds at every node of G. The n-flow is **nowhere zero** if $g(e) \neq 0$ for all $e \in E$.

(a) Show that if G has a nowhere-zero n-flow, then G has a nowhere-zero $\mathbb{Z}_n$-flow.

(b) For a nowhere-zero $\mathbb{Z}_n$-flow f, define $g : E \to [-(n-1), n-1]$ such that $g(e)$ is congruent to $f(e)$ modulo n. The conservation of flow of g is not necessarily satisfied at each node. The absolute value of the different between incoming and outgoing flow at v is called the **excess**.

An **augmenting path** from a node u to a node v is a path $u = u_0 u_1 \dots u_r = v$ in the undirected graph G such that $u_{i-1} \to u_i$ is a directed edge in ${}_\rho G$ if and only if $g(u_{i-1}u_i) > 0$. Let $h : E \to \{-1, 0, 1\}$ be the function such that $h(e)g(e) > 0$ if e is on the path and $h(e) = 0$ otherwise. Show that $g + nh : E \to \mathbb{Z}$ still takes values in the interval $[-(n-1), n-1]$ and reduces the excess at some node.

(c) Prove that if G has a nowhere-zero $\mathbb{Z}_n$-flow, then G has a nowhere-zero n-flow.

(d) Prove that

$$\varphi_G(n) \neq 0 \qquad \text{implies} \qquad \varphi_G(n+1) \neq 0\,.$$

(e) Even stronger, prove that

$$\varphi_G(n) \leq \varphi_G(n+1)\,.$$

(This is nontrivial. But you will easily prove this after having read Chapter 7.)

1.15 Let ${}_\rho G = (V, E, \rho)$ be an oriented graph and $n \geq 2$.

(a) Let $f : E \to \mathbb{Z}_n$ be a nowhere-zero $\mathbb{Z}_n$-flow and let $e \in E$. Show that f naturally yields a nowhere-zero $\mathbb{Z}_n$-flow on the contraction ${}_\rho G/e$.

(b) For $S \subseteq V$ let $E^{\text{in}}(S)$ be the **in-coming** edges, i.e., $u \to v$ with $v \in S$ and $u \in V \setminus S$, and let $E^{\text{out}}(S)$ be the **out-going** edges. Show that $f : E \to \mathbb{Z}_n$ is a nowhere-zero $\mathbb{Z}_n$-flow if and only if

$$\sum_{e \in E^{\text{in}}(S)} f(e) \;=\; \sum_{e \in E^{\text{out}}(S)} f(e)$$

for all $S \subseteq V$. (*Hint:* For the sufficiency, contract all edges in S and $V \setminus S$.)

(c) Infer that $\varphi_G \equiv 0$ if G has a bridge.

1.16 Discover the notion of tensions.

1.17 Consider the **Petersen graph** G pictured in Figure 1.17.

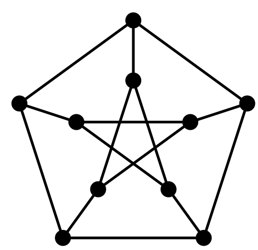

Figure 1.17. The Petersen graph.

(a) Show that $\varphi_G(4) = 0$.
(b) Show that the polynomial $\varphi_G(n)$ has nonreal roots.
(c) Construct a planar[4] graph whose flow polynomial has nonreal roots. (*Hint:* Think of the dual coloring question.)

1.18 ☖ Prove Proposition 1.3.3: Let ${}_\rho G = (V, E, \rho)$ be an acyclic graph and $\Pi = \Pi({}_\rho G)$ the induced poset. A map $c : V \to [n]$ is strictly compatible with the orientation ρ of G if and only if c is a strictly order-preserving map $\Pi \to [n]$.

1.19 Compute $\Omega^\circ_{D_{10}}(n)$.

1.20 ☖ Show that $\Omega_\Pi(n)$ is a polynomial in n.

1.21 ☖ Let $\mathcal{S} = \operatorname{conv}\left\{ \binom{a_1}{b_1}, \binom{a_2}{b_2} \right\}$, with $a_1, a_2, b_1, b_2 \in \mathbb{Z}$, be a **lattice segment**.[5] Show that

$$\operatorname{ehr}_\mathcal{S}(n) \;=\; L\,n + 1\,,$$

where $L = |\gcd(a_2 - a_1, b_2 - b_1)|$, the **lattice length** of $\mathcal{S}$. Conclude further that $-\operatorname{ehr}_\mathcal{S}(-n)$ equals the number of lattice points of $n\mathcal{S}$ other than the endpoints, in other words,

$$(-1)^{\dim \mathcal{S}} \operatorname{ehr}_\mathcal{S}(-n) \;=\; \operatorname{ehr}_{\mathcal{S}^\circ}(n)\,.$$

Can you find an explicit formula for $\operatorname{ehr}_\mathcal{S}(n)$ when $\mathcal{S}$ is a segment with *rational* endpoints?

[4]The Petersen graph is a (famous) example of a nonplanar graph.
[5]We use the notation $\operatorname{conv}(V)$ to denote the convex hull of a set V of vectors.

1.22 Let O be a closed polygonal lattice path, i.e., the union of lattice segments, such that each vertex on O lies on precisely two such segments, and that topologically O is a closed curve. Show that

$$\mathrm{ehr}_O(n) \ = \ L\,n\,,$$

where L is the sum of the lattice lengths of the lattice segments that make up O or, equivalently, the number of lattice points on O.

1.23 ☖ Let $\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{Z}^2$, and let Q be the half-open parallelogram

$$\mathsf{Q} := \{\lambda \mathbf{v}_1 + \mu \mathbf{v}_2 \,:\, 0 \le \lambda, \mu < 1\}\,.$$

Show (for example, by tiling the plane by translates of Q) that

$$\mathrm{ehr}_{\mathsf{Q}}(n) \ = \ A\,n^2,$$

where $A = \left|\det\left(\begin{smallmatrix} a & c \\ b & d \end{smallmatrix}\right)\right|$.

1.24 A lattice triangle $\mathrm{conv}\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$ is **unimodular** if $\mathbf{v}_2 - \mathbf{v}_1$ and $\mathbf{v}_3 - \mathbf{v}_1$ form a lattice basis of $\mathbb{Z}^2$.

(a) Prove that a lattice triangle is unimodular if and only if it has area $\frac{1}{2}$.

(b) Conclude that for any two unimodular triangles $\triangle_1$ and $\triangle_2$, there exist $T \in \mathrm{GL}_2(\mathbb{Z})$ and $\mathbf{x} \in \mathbb{Z}^2$ such that $\triangle_2 = T(\triangle_1) + \mathbf{x}$.

(c) Compute the Ehrhart polynomials of all unimodular triangles.

(d) Show that every lattice polygon can be triangulated into unimodular triangles.

(e) Use the above facts to give an alternative proof of Theorem 1.4.1.

1.25 Let $\mathsf{P} \subset \mathbb{R}^2$ be a lattice polygon, and denote the area of P by A, the number of integer points inside the polygon P by I, and the number of integer points on the boundary of P by B. Prove that

$$A \ = \ I + \tfrac{1}{2}B - 1$$

(a famous formula due to Georg Alexander Pick). Deduce from this formulas for the coefficients of the Ehrhart polynomial of P.

1.26 Let $\mathsf{P}, \mathsf{Q} \subset \mathbb{R}^2$ be lattice polygons, such that Q is contained in the interior of P. Generalize Exercise 1.25 (i.e., both a version of Pick's theorem and the accompanying Ehrhart polynomial) to the "polygon with a hole" $\mathsf{P} - \mathsf{Q}$. Generalize your formulas to a lattice polygon with n "holes" (instead of one).

1.27 Let $f(t) \ = \ a_d t^d + a_{d-1} t^{d-1} + \cdots + a_0 \in \mathbb{R}[t]$ be a polynomial such that $f(n)$ is an integer for every integer $n > 0$. Give a proof or a counterexample for the following statements:

(a) All coefficients a_j are integers.

(b) $f(n)$ is an integer for all $n \in \mathbb{Z}$.

(c) If $(-1)^k f(-n) \ge 0$ for all $n > 0$, then $k = \deg(f)$.

1.28 Suppose $f(t) = a_d t^d + a_{d-1} t^{d-1} + \cdots + a_0 \in \mathbb{R}[t]$ is a polynomial with $a_d > 0$. Prove that, if all roots of $f(t)$ have negative real parts, then each $a_j > 0$.