Felix Breuer

Ham Sandwiches, Staircases and Counting Polynomials

Ham Sandwiches, Staircases and Counting Polynomials

Dissertation vorgelegt von

Felix Breuer

Fachbereich für Mathematik und Informatik
Freie Universität Berlin

Berlin, 2009

Felix Breuer
Freie Universität Berlin
Institut für Mathematik
Arnimallee 3
D-14195 Berlin
Germany
felix.breuer@fu-berlin.de

Betreuer und erster Gutachter:

Prof. Dr. Martin Aigner
Freie Universität Berlin
Institut für Mathematik
Arnimallee 2
D-14195 Berlin
Germany
aigner@math.fu-berlin.de

Zweiter Gutachter:

Prof. Dr. Volkmar Welker
Philipps-Universität Marburg
Fachbereich Mathematik und Informatik
Hans-Meerwein-Straße
D-35032 Marburg
Germany
welker@mathematik.uni-marburg.de

Datum der Disputation: 17. November 2009

Umschlaggestaltung durch den Autor unter Verwendung von Inkscape und Blender
Satz durch den Autor unter Verwendung von LaTeX
Herstellung: Books on Demand GmbH, Norderstedt

Contents

Introduction

This thesis consists of four chapters that are largely independent. A brief outline of each of the four chapters is given below. A more elaborate introduction to each of the four topics is given at the beginning of each chapter.

Chapter 1. Uneven Splitting of Ham Sandwiches. Let $\mu_1, \ldots, \mu_n$ be continuous probability measures on $\mathbb{R}^n$ and $\alpha_1, \ldots, \alpha_n \in [0,1]$. When does there exist an oriented hyperplane H such that the positive half-space H^+ has $\mu_i(H^+) = \alpha_i$ for all $i \in [n]$? We call such a hyperplane an α-splitting. It is well known that α-splittings do not exist in general. The famous Ham Sandwich Theorem states that if $\alpha_i = \frac{1}{2}$ for all i, then α-splittings exist for any choice of the μ_i.

We give sufficient criteria for the existence of α-splittings for general $\alpha \in [0,1]^n$. To better keep track of the pairs (μ_i, α_i) we introduce auxiliary functions $f_1, \ldots, f_n : S^{n-1} \to \mathbb{R}^n$ with the property that for all i the unique hyperplane H_i with normal v that contains the point $f_i(v)$ has $\mu_i(H_i^+) = \alpha_i$. Our main result is that if $\operatorname{Im} f_1, \ldots, \operatorname{Im} f_n$ are bounded and can be separated by hyperplanes, then an α-splitting exists. Interestingly, the equivariant methods that are classically used to prove the Ham Sandwich Theorem and similar equipartition results cannot be easily applied to show the existence of uneven splittings. We present a novel approach based on the Poincaré-Miranda Theorem.

One important property of this result is that it can be applied even if the supports of the μ_i overlap. This gives a partial answer to a question of Stojmenović. The main result implies several other criteria as corollaries, the weakest of which was also obtained independently by Bárány et al. Also, it allows an easy corollary of the classical Ham Sandwich Theorem to be generalized.

Chapter 2. Staircases in $\mathbb{Z}^2$. *This part is joint work with Frederik von Heymann.* Motivated by the study of lattice points inside polytopes, we seek to understand the set of lattice points "close" to a rational line in the plane. To this end, we define a staircase in the plane to be the set of lattice point in the plane below a rational line that have Manhattan distance less than 1 to the line. This set of lattice points is closely related to the Beatty and Sturmian sequences defined in number theory, i.e. to sequences of the form $(\left\lfloor \frac{b}{a}(n-1) \right\rfloor - \left\lfloor \frac{b}{a}n \right\rfloor)_{n \in \mathbb{N}}$ for $a, b \in \mathbb{N}$ with $\gcd(a,b) = 1$. We present three characterizations of these sequences from a geometric point of view. One of these characterizations is known, two are new. The most important one is recursive and closely related to the Euclidean Algorithm. In particular, we obtain recursive descriptions of staircases, as well as of the sets of lattice points inside the fundamental parallelepipeds of rational cones and inside certain triangles in the plane.

We then present several applications of our geometric observations. 1) We give a new proof of Barvinok's Theorem in dimension 2. Barvinok's Theorem states that the generating function of the lattice points inside a rational simplicial cone can be written as a short rational function. While Barvinok uses a signed decomposition of the cone into unimodular cones, we use a partition of the cone into sets that have a short representation. 2) We give a recursion formula for Dedekind-Carlitz polynomials, i.e. polynomials of the form $\sum_{k=1}^{a-1} x^{k-1} y^{\left\lfloor \frac{b}{a}k \right\rfloor}$. This answers a question of Beck, Haase and Matthews. 3) We simplify Scarf's proof of White's Theorem, which characterizes lattice simplices that contain no lattice points except their vertices.

Chapter 3. Counting Functions and Reciprocity Theorems. *This part is joint work with Raman Sanyal.* Results that give the values of counting polynomials at negative integers a combinatorial interpretation are called reciprocity theorems. We consider five natural counting polynomials determined by graphs. The chromatic polynomial, the modular flow and tension polynomials and the integral flow and tension polynomials. Reciprocity theorems for all of these are known, except for the modular flow polynomial. We fill in this missing result, thus answering a question of Beck and Zaslavsky.

The key ingredient in our proof of the Modular Flow Reciprocity Theorem is our construction of a disjoint union of open polytopes whose Ehrhart function is the modular flow polynomial. Such a construction was not known before, and thus the method of Beck and Zaslavsky, who used inside-out Ehrhart-Macdonald Reciprocity to prove the Integral Flow Reciprocity Theorem, could not be applied. We use classical Ehrhart-Macdonald Reciprocity on each of the polytopes in the union to obtain our reciprocity result. In their unpublished manuscript [BB] Babson and Beck give a similar reciprocity result, independently from our work.

In the remainder of the chapter we relate our construction to the theory of inside-out polytopes and apply these methods to give a new proof of the Modular Tension Reciprocity Theorem. It turns out that these two reciprocity results give rise to a very nice interpretation of the Tutte polynomial as a counting polynomial which is already implicit in the work of Reiner [Rei99]. Moreover this approach provides a unified framework in which the value of the Tutte polynomial at *every* lattice point in the plane can be interpreted. The chapter is rounded off by direct induction proofs for both the reciprocity result and the interpretation of the Tutte polynomial.

Chapter 4. Counting Functions as Hilbert Functions. *This part is joint work with Aaron Dall.* Steingrímsson showed that the chromatic polynomial of a graph, shifted by one, is the Hilbert function of a relative Stanley-Reisner ideal. More precisely, given a graph G, Steingrímsson defines two square-free monomial ideals $I_1 \subset I_2 \subset \mathbb{K}[x_1, \ldots, x_n]$ such that the $k+1$-colorings of G are in bijection with the monomials of degree k inside I_2 but outside I_1.

Going beyond the shifted chromatic polynomial, the question arises, which of the five counting polynomials introduced in Chapter 3 are Hilbert functions of this type? All five are! Even the chromatic polynomial itself, without the shift by one. We show this by giving a general theorem, that provides a sufficient criterion for when the Ehrhart function of a relative polytopal complex is a Hilbert function of Steingrímsson's type. We give two proofs of this theorem, one from a more combinatorial and one from a more geometric point of view. Also we present variations of this result, where weaker hypotheses lead to weaker conclusions. When applying our criterion to the five counting polynomials we make use of some of the polytopal complexes constructed in Chapter 3.

A proof that the tension complex is homeomorphic to Steingrímsson's coloring complex, a combinatorial characterization of the face lattice of the tension polytope, and a characterization of the Hilbert functions of relative Stanley-Reisner ideals in terms of their coefficients complete the chapter.

Zusammenfassung

In meiner Dissertation behandele ich vier verschiedene Themen.

Das Ungleichmäßige Teilen von Ham Sandwiches. Seien stetige Wahrscheinlichkeitsmaße $\mu_1, \ldots, \mu_n$ auf $\mathbb{R}^n$ und reelle Zahlen $\alpha_1, \ldots, \alpha_n \in [0, 1]$ gegeben. Wann existiert eine orientierte Hyperebene H, so dass für den positiven Halbraum H^+ gilt $\mu_i(H^+) = \alpha_i$ für alle $i \in [n]$? Solch eine Hyperebene nennen wir eine α-Teilung. α-Teilungen existieren nicht im Allgemeinen. Das berühmte Ham Sandwich Theorem besagt, dass $(\frac{1}{2}, \ldots, \frac{1}{2})$-Teilungen immer existieren.

Wir geben ein hinreichendes Kriterium für die Existenz einer α-Teilung für allgemeines $\alpha \in [0, 1]^n$ an. Als Mittel zur Buchhaltung führen wir Hilfsfunktionen $f_1, \ldots, f_n : S^{n-1} \to \mathbb{R}^n$ ein, die die Eigenschaft haben, dass für alle i die eindeutige Hyperebene H mit Normale v, die den Punkt $f_i(v)$ enthält, die Eigenschaft $\mu_i(H^+) = \alpha_i$ hat. Unser Hauptergebnis ist nun: Falls die Mengen $\operatorname{Im} f_1, \ldots, \operatorname{Im} f_n$ beschränkt sind und sich durch Hyperebenen trennen lassen, dann existiert eine α-Teilung. Interessant dabei ist, dass die äquivarianten Methoden, die üblicherweise verwendet werden, um das Ham Sandwich Theorem zu beweisen, sich hier nicht ohne Weiteres anwenden lassen. Wir verwenden eine neue Methode, die auf dem Poincaré-Miranda Theorem basiert.

Eine wichtige Eigenschaft dieses Ergebnisses ist, dass es sich auch anwenden lässt, wenn sich die Träger der μ_i überlappen. Wir geben damit eine teilweise Antwort auf eine Frage von Stojmenović. Darüber hinaus erhalten wir mehrere Korollare. Eines dieser Korollare wurde unabhängig von meiner Arbeit von Bárány et al. entdeckt.

Treppen im Gitter $\mathbb{Z}^2$. *Gemeinsame Arbeit mit Frederik von Heymann.* Durch das Studium der Menge der Gitterpunkte in einem Polytop motiviert, versuchen wir die Menge der Gitterpunkte "nah an" einer rationalen Geraden in der Ebene zu verstehen. Eine Treppe ist die Menge der Gitterpunkte unterhalb einer Geraden die Manhattan-Abstand weniger als 1 zu der Geraden haben. Diese Menge von Gitterpunkten ist eng verwandt mit den aus der Zahlentheorie bekannten Beatty- und Sturm-Sequenzen rationaler Zahlen. Mittels unseres geometrischen Ansatzes erhalten wir drei Charakterisierungen dieser Sequenzen; eine davon ist bekannt, zwei sind neu. Die wichtigste Charakterisierung ist rekursiv und eng verwandt mit dem Euklidischen Algorithmus. Insbesondere erhalten wir rekursive Beschreibungen von Treppen sowie den Mengen von Gitterpunkten in Parallelepipeden und bestimmten Dreiecken.

Unsere geometrischen Überlegungen lassen sich auf verschiedene Weisen anwenden. Wir geben 1) einen neuen Beweis von Barvinok's Theorem in Dimension 2, 2) eine Rekursionsformel für Dedekind-Carlitz Polynome und 3) einen teilweise neuen Beweis von Whites Charakterisierung leerer Gitterpolytope.

Zählpolynome und Reziprozitätssätze. *Gemeinsame Arbeit mit Raman Sanyal.* Reziprozitätssätze sind Theoreme, die dem Wert eines Zählpolynoms an einer negativen Zahl eine kombinatorische Interpretation geben. Wir betrachten fünf natürliche durch einen Graphen gegebene Zählpolynome: das chromatische Polynom, die k- und $\mathbb{Z}_k$-Flusspolynome und die k- und $\mathbb{Z}_k$-Tensionpolynome. Für alle diese sind Reziprozitätssätze bekannt, außer für das $\mathbb{Z}_k$-Flusspolynom. Wir ergänzen dieses fehlende Ergebnis und beantworten damit eine Frage von Beck und Zaslavsky. Ein ähnliches Ergebnis findet sich in einem unveröffentlichten Manuskript von Babson und Beck. Wir erhalten unser Ergebnis, indem wir das $\mathbb{Z}_k$-Flusspolynom als Ehrhart-Funktion einer disjunkten Vereinigung von Polytopen auffassen, aber wir stellen auch den Zusammenhang zur Theorie der Inside-Out Polytope her.

Die Reziprozitätssätze für die $\mathbb{Z}_k$-Fluss- und Tensionpolynome zusammen führen zu einer Interpretation des Tutte-Polynoms als Zählfunktion, die implizit schon bei Reiner zu finden ist. Darüber hinaus führt dieser Ansatz zu einer allgemeinen Methode, mit der die Werte des Tutte-Polynoms an allen Gitterpunkten in der Ebene interpretiert werden können. Schließlich geben wir noch direkte Induktionsbeweise für unseren Reziprozitätssatz und die Tutte-Interpretation.

Zählpolynome als Hilbert-Funktionen. *Gemeinsame Arbeit mit Aaron Dall.* Steingrímsson hat gezeigt, dass das chromatische Polynom eines Graphen, um eins verschoben, die Hilbert-Funktion eines relativen Stanley-Reisner-Ideals ist. Wie sieht es mit den fünf oben genannten Zählfunktionen aus? Wir zeigen, dass alle fünf diese Eigenschaft haben. Sogar das chromatische Polynom selbst, ohne Verschiebung. Zu diesem Zweck geben wir ein hinreichendes Kriterium dafür an, wann die Ehrhart-Funktion eines relativen polytopalen Komplexes die Hilbert-Funktion eines relativen Stanley-Reisner-Ideals ist. Dieses Kriterium beweisen wir auf zwei verschiede Weisen, geometrisch und kombinatorisch. Auch geben wir verschiedene Varianten dieses Kriteriums an, in denen abgeschwächte Voraussetzungen zu schwächeren Ergebnissen führen.

Dieser Abschnitt wird abgerundet durch 1) einen Beweis, dass der Tension-Komplex homöomorph zu Steingrímsson's Färbungskomplex ist, 2) eine kombinatorische Charakterisierung des Seitenverbands des Tension-Polytops und 3) eine Charakterisierung der Koeffizienten der Hilbert-Funktionen von relativen Stanley-Reisner-Idealen.

Acknowledgments

I want to thank my advisor Martin Aigner for his support and for the complete freedom he allowed me in deciding what I want to work on. I also thank him for his encouragement, back in 2006, to apply for a scholarship at the research training group "Methods for Discrete Structures" (GRK 1408). I am grateful to the Deutsche Forschungsgemeinschaft for supporting me during the three years in which all of this research was done and I am grateful to everyone at the MDS: for organizing such a wealth of mathematical activities and for making the research training group a great place to find collaborators to work with.

This brings me to my coauthors, to whom I am greatly indebted. Doing research together is a great experience with many merits of which I want to name but one: It makes the many frustrations inherent in mathematical research so much easier to bear.

I thank Frederik von Heymann for investigating staircases with me – or extra-coffee sequences as we used to call them. When we started out with the project, it was not clear at all whether the question we considered was suitable for a research project. I thank Fred for his interest, his enthusiasm and his courage to pursue the idea nonetheless.

I thank Aaron Dall for sharing the summer of 2008 and the research question he was currently working on with me. Thanks to him I learned to appreciate Hilbert functions and my three months in Berkeley were not only a great time but also a mathematical success.

I thank Raman Sanyal for his relentless energy. When I came back from Berkeley we started out trying to generalize the work Aaron and I had done and ended up proving a reciprocity theorem. I also thank him, Vanessa and Konstantin for inviting me in May 2009.

I thank Paul Bonsma for working with me on the recognition of fullerene patches. Even though the results of our project are not part of this thesis, I learned a lot during our long discussions about how best to write down our proofs.

The best part of my time as a PhD student were my three months in the Bay Area. For this I want to thank Matthias Beck and Tendai Chitewere, whose hospitality, kindness and friendship made it possible. I also want to thank Matthias Beck for his enthusiasm about mathematics, his interest in my work and his encouragement. His positive feedback was invaluable to me.

I thank both Christian Haase and Matthias Beck for suggesting that Aaron and I work on capturing the tension polynomial as a Hilbert function and for making the Doc-Course in the summer of 2007 such a stimulating event. I also thank Christian Haase for drawing the attention of Aaron and me to the concept of pulling triangulations and the attention of Fred and me to White's Theorem. And I am grateful to him, Andreas Paffenholz and Benjamin Nill for running the research group lattice polytopes at the FU.

I thank Mark de Longueville for getting me interested in topological combinatorics when I was still a graduate student and for several helpful conversations about ham sandwiches and other things.

There are a number of other people whose comments have helped me along the way. I thank Stefan Felsner for coming up with a counterexample, when I was still working on a different topic and I thank the referees of the journal Discrete and Computational Geometry for their comments on my paper [Bre09]. I thank Günter Rote pointing me to [Sto90], Jesús Jerónimo for pointing me to [BHJ08, Kin08] and Victor Alvarez for pointing us to [BGMM99]. And I thank Kathrin Vorwerk, Holger Dell and Jiří Fink with whom I puzzled over various problems, unrelated to this thesis.

My warmest thanks go to my friends and family. I am infinitely grateful for all those conversations: in offices or on park benches, in Berkeley, Magdeburg or Seville, about math, politics, technology, humor, motivation, work, love and life. Here I want to mention Peter Krautzberger, Thomas Friedrich, Cornelia Dangelmayr, Mareike Massow, Daria Schymura, Julia Westendorf, Thomas Berndt, Ellen Hüther and most importantly Doris Breuer and Chambui Jugder. Without Doris I wouldn't be who I am today and thanks to Chambui I did not only brave the ups and downs that came with writing this dissertation, but thanks to her the last years were the best I have had.

Chapter 1.

Uneven Splitting of Ham Sandwiches

The Ham Sandwich Theorem states that any ham sandwich in $\mathbb{R}^3$ consisting of ham, cheese and bread can be split by a hyperplane such that ham, cheese and bread are simultaneously split in half. In general, it reads as follows.

1.0.1. Ham Sandwich Theorem.

Let $\mu_1, \ldots, \mu_n$ be continuous probability measures on $\mathbb{R}^n$. Then there exists an oriented hyperplane H such that

$$\mu_i(H^+) = \frac{1}{2} \text{ for all } i \in \{1, \ldots, n\} =: [n],$$

where H^+ denotes the positive half-space corresponding to H.

The Ham Sandwich Theorem was posed as a problem by Steinhaus in the Scottish Book [Mau81, Problem 123]. A proof for $n = 3$ was given by Banach [Ste38] using the Borsuk-Ulam Theorem. This proof was generalized to arbitrary n by Stone and Tukey in [ST42]. For an account of the early history of the Ham Sandwich Theorem we refer to [BZ04]. For a reference on the Borsuk-Ulam Theorem, including its application to the Ham Sandwich Theorem and an overview of related results, we recommend [Mat08].

The question we want to address is this: Let $\mu_1, \ldots, \mu_n$ be continuous probability measures on $\mathbb{R}^n$ and $\alpha_1, \ldots, \alpha_n \in [0, 1]$. When does there exist a hyperplane H such that $\mu_i(H^+) = \alpha_i$ for all $i \in [n]$?

It is well-known that such a hyperplane does not exist in general [ST42], [Mat08]. We give examples for this in Section 1.4. Interestingly the equivariant methods used to obtain equipartition results cannot be easily applied to obtain partitions according to other ratios.

The purpose of this chapter is to give a sufficient criterion for the existence of a hyperplane splitting continuous probability measures $\mu_1, \ldots, \mu_n$ in $\mathbb{R}^n$ according to prescribed ratios $\alpha_1, \ldots, \alpha_n \in [0,1]$. The idea is this: Given a ratio $\alpha_i \in [0,1]$ and a measure μ_i we define an auxiliary function f_i that selects for any normal vector $v \in S^{n-1}$ a point $f_i(v) \in H$ from a hyperplane H with $\mu_i(H^+) = \alpha_i$. We do not assume the continuity of the f_i, only that their image Im f_i is bounded. Our main result is: if the sets Im $f_1, \ldots,$ Im f_n can be separated by hyperplanes, then there exists a hyperplane H with $\mu_i(H^+) = \alpha_i$ for all i. We call such a hyperplane an α-splitting. One corollary is that if the supports of the measures μ_i can be separated, we can find a splitting hyperplane for *any* ratios α_i, but it is important to stress that the main result can also be applied if the supports of the μ_i overlap. Thus we give a partial answer to a question posed in [Sto91]. Interesting about these results is that we are able to show the existence of the desired hyperplane despite the absence of symmetry that is usually afforded by the ratio $\frac{1}{2}$. To this end we introduce a new method based on the Poincaré-Miranda Theorem.

There is much literature on the Ham Sandwich Theorem and related partitioning results, see [Mat08] for an overview and e.g. [BM01], [BKS00], [Dol92], [Grü60], [GH05], [Rad46], [Ram96], [VŽ01], [ŽV90]. However, most results assert the existence of equipartitions. Some sources in which the existence of partitions into parts of different measure are shown, are the following.

A k-fan in $\mathbb{R}^2$ is a set of k rays emanating from a point p. The regions in between two adjacent rays are called sectors $\sigma_1, \ldots, \sigma_k$. In the case of continuous measures a k-fan $(\alpha_1, \ldots, \alpha_k)$-partitions measures $\mu_1, \ldots, \mu_n$ if $\mu_i(\sigma_j) = \alpha_j$ for all j and i. Bárány and Matoušek [BM01] show that any two measures on $\mathbb{R}^2$ can be α-partitioned by a 2-fan for all α and by a 4-fan for $\alpha = (\frac{2}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5})$.

Vrećica and Živaljević [VŽ01] show that for any continuous measure, any α with $\sum_i \alpha_i = 1$, any non-degenerate simplex Δ in $\mathbb{R}^n$ and any point $a \in \operatorname{int} \Delta$ there is a vector $v \in \mathbb{R}^n$ such that the cones R_i with apex a generated by the facets of Δ satisfy $\mu(R_i + v) = \alpha_i$ for all i.

Rado [Rad46] shows that for any measure μ on $\mathbb{R}^n$ there exists a point x such that $\mu(H^+) \geq \frac{1}{n+1}$ for any half-space H^+ with $x \in H^+$. Živaljević and Vrećica [ŽV90] give a unification of Rado's Theorem with the Ham Sandwich Theorem: For measures $\mu_1, \ldots, \mu_k$ on $\mathbb{R}^n$ there is a $(k-1)$-dimensional affine subspace A such that $\mu(H^+) \geq \frac{1}{n-k+2}$ for any half-space H^+ with $A \subset H^+$. See also [Dol92],[Vre03].

Stojmenović [Sto90] gives an algorithm for finding a hyperplane bisecting the volume of three convex polygons in $\mathbb{R}^3$ that can be separated by hyperplanes (see Section 1.5). He remarks that this algorithm works in any dimension, for other measures (not only volume) and other proportions of splitting, thus implicitly asserting the existence of a splitting in these cases. This assertion amounts to Corollary 1.9.4. In the journal version [Sto91] of that article, however, Stojmenović only remarks that extensions to other measures and proportions are possible, but no mention is made of dimensions $n > 3$. Stojmenović also asked for a criterion that can be applied in the case that the supports are not disjoint; Theorem 1.6.1 and Corollary 1.9.2 give a partial answers to this question.

After the contents of this chapter were published online [Bre09], Jesús Jerónimo-Castro informed the author of the articles [BHJ08] and [Kin08]. In [BHJ08] Bárány, Hubard and Jerónimo give a proof of Corollary 1.9.4 in the special case that the measures considered have an additional "nice" property. Their proof is different from mine. See Section 1.9. In [Kin08] Kincses shows that if we are given k convex bodies in $\mathbb{R}^n$ that can be separated by hyperplanes and the measures $\mu_1, \ldots, \mu_k$ are given by their respective volumes, then the space of all α-splittings is diffeomorphic to the sphere S^{n-k}.

The organization of this chapter is as follows. In Section 1.1 we give preliminary definitions that are also of relevance for the following chapters, before we state the problem we discuss in this chapter formally in Section 1.2. In Section 1.3 we introduce the auxiliary functions f_i. Examples showing the non-existence of uneven splittings in general are given in Section 1.4, which motivate the concept of separability defined in Section 1.5. We then have everything in place to state our main result (Theorem 1.6.1) in Section 1.6. Before turning to the proof in Section 1.8 we take a look at the main tool, the Poincaré-Miranda Theorem, in Section 1.7. There are many ways of choosing the auxiliary functions f_i and each gives rise to a different corollary of Theorem 1.6.1. For the purposes of Theorem 1.6.1 only the sets Im f_i are of interest. In Section 1.9 we give a method of guaranteeing the existence of functions f_i while controlling their image. This gives rise to Corollaries 1.9.2 and 1.9.4. The latter is the aforementioned result that guarantees the existence of α-splittings if the supports can be separated, while the former is a stronger statement that can also be applied if the supports overlap. We apply this method in Section 1.10 to generalize an easy corollary of the classical Ham Sandwich Theorem on the partitioning of one mass by two lines. We conclude the chapter in Section 1.11 by suggesting a canonical choice for the functions f_i.

1.1. Convex Geometry

In this section we gather some preliminary definitions and notational conventions, mainly from convex geometry. For more detail we recommend [Zie97], [Sch86], [Mat08] and [Hat02].

The **convex hull** of a set S is denoted by $\mathsf{conv}(S)$ and the **affine hull** is denoted by $\mathsf{aff}(S)$.

A **hyperplane** in $\mathbb{R}^n$ is an affine subspace of dimension $n-1$. Any hyperplane H separates $\mathbb{R}^n$. The two connected components of $\mathbb{R}^n \setminus H$ are the **open half-spaces** determined by H. Their respective closures are the **closed half-spaces**. If we just speak of a half-space, we always refer to the respective closed half-space. A hyperplane is **oriented** if one of the two induced hyperplanes is labeled **positive** while the other is labeled as **negative**. The positive half-space is denoted H^+ and the negative half-space is denoted H^-. Thus the open positive half-space is $\mathsf{int}\, H^+$ while the open negative half-space is $\mathsf{int}\, H^-$.

We will frequently encounter such labelings throughout all chapters of this thesis. We will always use $+$ as the label of the positive part and $-$ as the label of the negative part. In any such context we identify $+$ with the number $+1$ and $-$ with the number -1. Thus in the case of half-spaces $H^+ = H^{+1}$ and $H^- = H^{-1}$.

An orientation of a hyperplane H is determined by fixing a normal vector v of H. The half-space in the direction of v is positive, while the half-space in the direction of $-v$ is negative.

All positive multiples of v lead to the same orientation of H, while the negative multiples of v lead to the opposite orientation. Thus we can parameterize the hyperplanes and half-spaces in $\mathbb{R}^n$ by pairs (v, λ) of a normal vector $v \in \mathbb{R}^n \setminus \{0\}$ and a scalar $\lambda \in \mathbb{R}$ in the following fashion:

$$\begin{aligned} H_{v,\lambda} &= \{x \in \mathbb{R}^n : \langle x, v\rangle = \lambda\}, \\ H^+_{v,\lambda} &= \{x \in \mathbb{R}^n : \langle x, v\rangle \geq \lambda\}, \\ H^-_{v,\lambda} &= \{x \in \mathbb{R}^n : \langle x, v\rangle \leq \lambda\}. \end{aligned}$$

A **polyhedron** P in $\mathbb{R}^n$ is the intersection of finitely many closed half-spaces. An **open** polyhedron is the intersection of finitely many open half-spaces with an affine space. A **polytope** is a bounded polyhedron. Equivalently a polytope is the convex hull of finitely many points in $\mathbb{R}^n$. The **dimension** $\mathsf{dim}\, P$ of a polyhedron P is the smallest dimension of an affine subspace containing P. The proper **faces** of a polyhedron P are the polytopes of the form $H_{v,\lambda} \cap P$ where $H^+_{v,\lambda} \supset P$. Additionally, P is a face of itself. We write $F \leq P$ to denote that F is a face of P. The **facets** of P are the faces F of P with $\mathsf{dim}\, F = \mathsf{dim}\, P - 1$. The **vertices** of P are the 0-dimensional faces of P. A vertex is a singleton set $\{v\} \subset \mathbb{R}^n$ and we will also call its one element v a vertex. In general we will not distinguish, throughout this thesis, between a singleton set and its one element. We denote the set of vertices of a polyhedron P with $\mathsf{vert}(P)$. A d-dimensional **simplex** is the convex hull of $d + 1$ affinely independent points. All faces of a simplex are simplices.

A **cone** C in $\mathbb{R}^n$ is a convex set that is closed under taking non-negative multiples, i.e. C is a cone C is convex and if $x \in C$ implies $\lambda x \in C$ for any $0 \leq \lambda \in \mathbb{R}$. A **polyhedral cone** is a cone that is also a polyhedron. *All cones in this thesis will be polyhedral*, so we henceforth make no mention of the adjective “polyhedral”. A cone is **pointed** if it does not contain a linear subspace as a subset. $\mathsf{cone}(S)$ denotes the cone generated by a set S, that is the minimal cone containing S. If P is a polytope then $\mathsf{cone}(P)$ is polyhedral.

A polyhedron P is **rational** if P can be written as the intersection of half-spaces $H^+_{v,\lambda}$ such that all entries of v and λ are rational. For polytopes P it is equivalent to say that P is rational if P can be written as the convex hull of finitely many points with rational coordinates. In *this* chapter we will not be dealing with rational polyhedra, but with polyhedra in general. However, *in all other chapters, all polyhedra will be rational*. We will not mention the adjective “rational” in the following chapters; by convention all polyhedra in chapters 2, 3 and 4 are rational.

A **polyhedral complex** $\mathcal{C}$ is a finite set of polyhedra such that if $F \leq P \in \mathcal{C}$ then $F \in \mathcal{C}$ and if $P_1, P_2 \in \mathcal{C}$ then $P_1 \cap P_2 \in \mathcal{C}$ and $P_1 \cap P_2$ is a face of both P_1 and P_2. The elements of $\mathcal{C}$ are also called **faces**. The **dimension** of a polyhedral complex is the maximal dimension of a face of $\mathcal{C}$. The **facets** are the inclusion-maximal faces of $\mathcal{C}$. A polyhedron together with all its faces also constitutes a polyhedral complex and we will frequently view a polyhedron as such. Note that the facets of a polyhedron P are the codimension 1 faces, while the only facet of P viewed as a polyhedral complex is P itself. It will be clear from context which “facet” is meant. The **vertices** of $\mathcal{C}$ are its 0-dimensional faces; the set of vertices is denoted by $\mathsf{vert}(\mathcal{C})$. A **polytopal complex** is a polyhedral complex consisting of polytopes and a (geometric)

simplicial complex is a polyhedral complex consisting of simplices. A **fan** is a polyhedral complex whose elements are cones. The **face fan** of polytope P is the fan consisting of the cones $\mathsf{cone}(F)$ for all proper faces F of P. The notation $\bigcup \mathcal{C}$ denotes the union of all faces of $\mathcal{C}$; $\bigcup \mathcal{C}$ is simply the set of points in $\mathbb{R}^n$ that are contained in any face of $\mathcal{C}$. A **subdivision** of a polyhedral complex $\mathcal{C}$ is a polyhedral complex $\mathcal{C}'$ with $\bigcup \mathcal{C} = \bigcup \mathcal{C}'$ such that every face of $\mathcal{C}'$ is contained in a face of $\mathcal{C}$. A triangulation is a subdivision which is simplicial.

The **relative interior** $\mathsf{relint}\, P$ of a polyhedron P is the interior of P taken with respect to the affine hull $\mathsf{aff}(P)$. Similarly, the **boundary** ∂P is the boundary of P taken with respect to the affine hull $\mathsf{aff}(P)$. $\mathsf{relint}\, P$ and ∂P are disjoint. The proper faces of P are contained in ∂P. Every $x \in \bigcup \mathcal{C}$ is contained in a unique inclusion-minimal face of $\mathcal{C}$. Equivalently, every $x \in \bigcup \mathcal{C}$ is contained in the relative interior of a unique face of $\mathcal{C}$. This unique face is denoted by $\mathsf{minface}_{\mathcal{C}}(x)$.

A **hyperplane arrangement** is a finite set of hyperplanes $\mathcal{H}$ in some $\mathbb{R}^n$. $\mathcal{H}$ is called **central** if all hyperplanes in $\mathcal{H}$ are linear. A hyperplane arrangement $\mathcal{H}$ defines a polyhedral complex $\mathsf{comp}(\mathcal{H})$: Let $S := \{H^+ : H \in \mathcal{H}\} \cup \{H^- : H \in \mathcal{H}\}$ be the set of all half-spaces given by hyperplanes in $\mathcal{H}$. Then the faces of $\mathsf{comp}(\mathcal{H})$ are all intersections of half-spaces in S. If $\mathcal{H}$ is central, then $\mathsf{comp}(\mathcal{H})$ is a fan.

Given two polyhedral complexes $\mathcal{C}_1$ and $\mathcal{C}_2$, their **intersection** is the polyhedral complex $\mathcal{C}_1 \cap \mathcal{C}_2 := \{P_1 \cap P_2 : P_1 \in \mathcal{C}_1, P_2 \in \mathcal{C}_2\}$. This is obviously not the usual set theoretic notion of intersection. With this definition $\bigcup(\mathcal{C}_1 \cap \mathcal{C}_2) = \bigcup \mathcal{C}_1 \cap \bigcup \mathcal{C}_2$.

An **abstract** or **combinatorial simplicial complex** is a collection Δ of subsets of a given finite ground set S such that $\sigma_1 \subset \sigma_2 \in \Delta$ implies $\sigma_1 \in \Delta$. The vertices of Δ are the 1 element sets in Δ; the vertex set of Δ is denoted by $\mathsf{vert}(\Delta)$. If $\mathcal{C}$ geometric simplicial complex then the set $\mathsf{comb}(\mathcal{C}) := \{\mathsf{vert}(P) : P \in \mathcal{C}\}$ is a combinatorial simplicial complex on the ground set $\mathsf{vert}(\mathcal{C})$. Note that in general the ground set of an abstract simplicial complex and its vertex set need not coincide. This fine point will only be relevant in Chapter 4. Note that every abstract simplicial complex can be realized as a geometric simplicial complex in some $\mathbb{R}^d$.

1.2. The Problem

By a **continuous probability measure** μ on $\mathbb{R}^n$ we mean a measure μ on the set $\mathbb{R}^n$ equipped with the Borel σ-algebra such that $\mu(\mathbb{R}^n) = 1$ and μ is absolutely continuous with respect to the Lebesgue measure L, i.e. if $L(S) = 0$, then $\mu(S) = 0$ for all measurable sets S. Whenever we use the term "measure", we assume that all of these conditions hold. Note that for a continuous measure the map $(v, \lambda) \mapsto \mu(H^+_{v,\lambda})$ is continuous. Note also that any continuous measure μ can be written as $\mu(S) = \int_S h\, dL$ where h is a Borel-measurable function by the Radon-Nikodym Theorem. Two h_1, h_2 with this property have to agree almost everywhere. We think of μ as coming with a fixed choice of h and in abuse of terminology define the **support** of a measure as the set where h is non-zero.

The question we want to address is the following.

1.2.1. Question.

Let $\mu_1, \ldots, \mu_n$ be continuous probability measures on $\mathbb{R}^n$ and $\alpha_1, \ldots, \alpha_n \in [0,1]$. When does there exist a hyperplane $H_{v,\lambda}$ such that $\mu_i(H^+_{v,\lambda}) = \alpha_i$ for all $i \in [n]$?

We call such a hyperplane $H_{v,\lambda}$ an $(\alpha_1, \ldots, \alpha_n)$**-splitting** of $\mu_1, \ldots, \mu_n$. If $\alpha_i \neq \frac{1}{2}$ for some i, we call the splitting **uneven**. It is well-known that for a given uneven α, α-splittings do not exist in general. Two examples for this are presented in Section 1.4.

1.3. *Auxiliary functions f_i*

Before we turn to these examples we introduce a tool to keep track of the set of hyperplanes H with $\mu(H) = \alpha$, for a fixed measure μ and a fixed ratio $\alpha \in [0,1]$. Unless the support of μ is unbounded and $\alpha \in \{0,1\}$, we always have that for any $v \in \mathbb{R}^n \setminus \{0\}$ there exists a hyperplane $H_{v,\lambda}$ with $\mu(H^+_{v,\lambda}) = \alpha$. The idea is to pick a point $f(v) \in H_{v,\lambda}$ to record its position. More formally, we let $f : \mathbb{R}^n \setminus \{0\} \to \mathbb{R}^n$ be a function with $\mu(H^+_{v,\langle v,f(v)\rangle}) = \alpha$ for all v. Such an f we call an **auxiliary function**. Note that $H_{v,\langle v,x\rangle}$ is the unique hyperplane with normal v that contains x. Note also that in general, we do not require f to be continuous. We will sometimes view f as defined only on S^{n-1}.

1.3.1. Example. As an example consider the probability measure μ defined by the probability density that is constant non-zero on a ball $B \subset \mathbb{R}^n$ with center c and radius r and zero everywhere else. Let $0 < \alpha < \frac{1}{2}$. In this case we are going to construct a particular auxiliary function f as follows. See Figure 1.1. First we note that for any $v \in \mathbb{R}^n \setminus \{0\}$ there is a unique $\lambda \in \mathbb{R}$ such that $\mu(H^+_{v,\lambda}) = \alpha$. We chose $f(v)$ to be the center of the disk $H_{v,\lambda} \cap B$. Then $f : S^n \to \mathbb{R}^n$ is continuous and differentiable and Im f is a sphere with center c and radius strictly smaller than r. Moreover for a given v the unique hyperplane $H_{v,\lambda}$ such that $\mu(H^+_{v,\lambda}) = \alpha$ is tangent to Im f with $H^-_{v,\lambda} \supset \text{Im } f$. These additional properties of f aid intuition, however, we make use of them only in the examples in Section 1.4. We will say more on this particular way of constructing an auxiliary function in Section 1.11.

To illustrate the use of such a map f, we give the following proposition.

1.3.2. Proposition.

Let $\mu_1, \ldots, \mu_n$ be continuous measures, $\alpha_1, \ldots, \alpha_n \in (0,1)$ such that for all i and all $v \in S^n$ the choice of λ such that $\mu_i(H^+_{v,\lambda}) = \alpha_i$ is unique. Let $f_i : S^n \to \mathbb{R}^n$ be such that $\mu_i(H^+_{v,\langle v,f(v)\rangle}) = \alpha_i$ for all i.

Then there is a hyperplane H such that $\mu_i(H^+) = \alpha_i$ for all i if and only if there exists a $v \in S^n$ and a hyperplane H with normal v such that $f_i(v) \in H$ for all i.

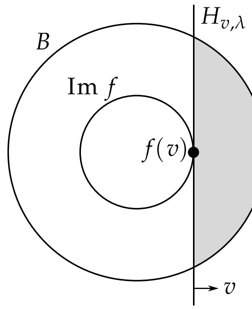

Figure 1.1: Construction of an auxiliary function for the measure μ whose density is constant non-zero on a disk B and zero everywhere else. For a given normal v the hyperplane $H_{v,\lambda}$ is such that the shaded area has measure $\mu(H^+_{v,\lambda}) = \alpha$. We need to pick one point $f(v) \in H_{v,\lambda}$. In this example we chose $f(v)$ to be the center point of the line $B \cap H_{v,\lambda}$ for all v. Then Im f is a circle.

Proof. "$\Rightarrow$" Let v be the normal of H. By assumption there is only one hyperplane with normal v that gives the desired split. So all $f_i(v)$ have to be contained in H.

"$\Leftarrow$" If there are a vector v and a hyperplane H with normal v such that $f_i(v) \in H$ for all i, then $H = H_{v,\langle v, f_i(v)\rangle}$ for all i and hence $\mu_i(H^+) = \alpha_i$ for all i as desired. □

1.4. Sandwiches Without Uneven Splittings

We are now going to take an informal look at two examples where there are no α-splittings, for certain uneven α. These examples are built from spheres as in Example 1.3.1. Note that in this case, for $\alpha \in (0,1)$ the normal vector determines the hyperplane H with $\mu(H^+) = \alpha$ uniquely.

1.4.1. Example. Let μ_1 and μ_2 be two measures in $\mathbb{R}^2$ given by two concentric discs of different radius as in Example 1.3.1. Say the radius of the disc of μ_1 is larger. Fix $\alpha_1 = \alpha_2 = \alpha \in (0, \frac{1}{2})$ and consider the corresponding f_1, f_2 as defined in Example 1.3.1. Both Im f_1 and Im f_2 are circles with same center and different radii, the former containing the latter. Suppose there is a hyperplane $H_{v,\lambda}$ such that $\mu_1(H^+_{v,\lambda}) = \mu_2(H^+_{v,\lambda}) = \alpha$. Then $H^+_{v,\lambda}$ would have to be tangent to both circles Im f_1 and Im f_2, which is impossible. Hence, there is no α-splitting.

1.4.2. Example. Let μ_1, μ_2, μ_3 be given by 3 disjoint balls in $\mathbb{R}^3$ such that their center-points lie on a line and the radius of the center sphere (say μ_2) is larger than that of the other two. Fix one $\alpha \in (0, \frac{1}{2})$ for all μ_i. We construct the auxiliary functions as in Example 1.3.1. The resulting f_i are again spheres and the center one has larger radius than the other two. See

Figure 1.2. Now any hyperplane tangent to Im f_1 and Im f_2 cannot meet Im f_3, so there is no hyperplane splitting all three measures according to the desired ratio.

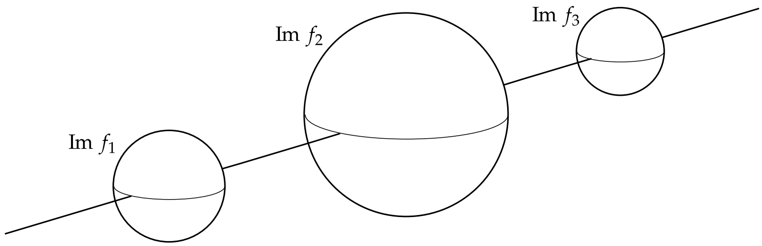

Figure 1.2: The images of three auxiliary functions constructed as in Example 1.3.1. An α-splitting would be a hyperplane tangent to all three spheres, but such a hyperplane does not exist.

1.5. *Separability*

We have seen that if the Im f_i lie on a line, we cannot expect there to be an uneven splitting. However if the Im f_i were to be in "general position", we might have more luck. To make this notion precise we define the concept of sets in n-space that can be separated by hyperplanes.

n points in $\mathbb{R}^n$ are affinely independent, if and only if for any partition of these points into two classes, there is an affine hyperplane such that the one class is on one side and the other class is on the other side. In analogy, we say that sets $S_1, \ldots, S_n \subset \mathbb{R}^n$ can be strictly separated by hyperplanes, or **separated** for short, if for any function $\sigma : [n] \to \{-1, +1\}$, there is a hyperplane H such that $\operatorname{int} H^{\sigma(i)} \supset S_i$ for all i. In other words, no matter how we prescribe which S_i are supposed to be in front of H and which S_i are supposed to be behind H, we can always find a hyperplane H that achieves this. See Figure 1.3 for an example.

To illustrate this definition we give the following proposition.

1.5.1. Proposition.

1. Sets $S_1, \ldots, S_n \subset \mathbb{R}^n$ can be separated if and only if $\operatorname{conv} S_1, \ldots, \operatorname{conv} S_n$ can be separated.
2. Convex sets $S_1, \ldots, S_n \subset \mathbb{R}^n$ can be separated if and only if there does not exist an $(n-2)$-dimensional affine subspace that meets all of the S_i.
3. Points $s_1, \ldots, s_n \in \mathbb{R}^n$ are in general position if and only if the singletons $\{s_1\}, \ldots, \{s_n\} \subset \mathbb{R}^n$ can be separated.

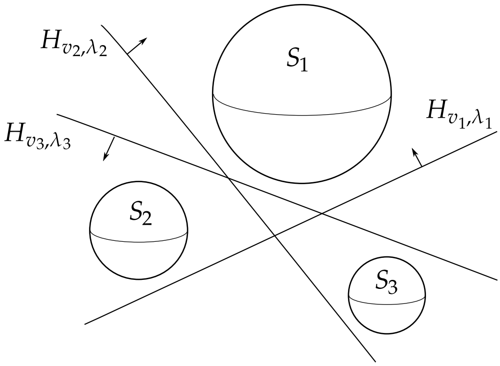

	S_1	S_2	S_3
H_{v_4,λ_4}	$+$	$+$	$+$
H_{v_1,λ_1}	$+$	$+$	$-$
H_{v_2,λ_2}	$+$	$-$	$+$
H_{v_3,λ_3}	$-$	$+$	$+$
$H_{-v_1,-\lambda_1}$	$-$	$-$	$+$
$H_{-v_2,-\lambda_2}$	$-$	$+$	$-$
$H_{-v_3,-\lambda_3}$	$+$	$-$	$-$
$H_{-v_4,-\lambda_4}$	$-$	$-$	$-$

Figure 1.3: The figure shows three spheres $S_1, S_2, S_3 \subset \mathbb{R}^3$ and three lines which define hyperplanes $H_{v_1,\lambda_1}, H_{v_2,\lambda_2}, H_{v_3,\lambda_3} \subset \mathbb{R}^3$ that are orthogonal to the plane of projection. Let H_{v_4,λ_4} denote a fourth hyperplane such that $S_1, S_2, S_3 \subset H^+_{v_4,\lambda_4}$. The table on the right-hand side lists for each hyperplane (and for the hyperplanes with opposite orientation) which spheres are contained in the positive and negative half-space, respectively. All sign patterns σ appear. Hence the spheres can be separated.

1. is immediate from the definition. 2. can be found in e.g. [Mat02, p. 218]. 3. is an immediate consequence of 2.

1.6. *The Main Result*

1.6.1. Theorem.

Let $\mu_1, \ldots, \mu_n$ be continuous probability measures on $\mathbb{R}^n$. Let $\alpha_1, \ldots, \alpha_n \in [0, 1]$. Let $f_1, \ldots, f_n : S^{n-1} \to \mathbb{R}^n$ be functions such that $\mu_i(H^+_{v,\langle v, f_i(v)\rangle}) = \alpha_i$ for all $i \in [n]$ and all $v \in S^{n-1}$. The f_i need not be continuous, but the $\operatorname{Im} f_i$ have to be bounded.

If $\operatorname{Im} f_1, \ldots, \operatorname{Im} f_n$ can be separated by hyperplanes, then there exists a hyperplane H such that $\mu_i(H^+) = \alpha_i$ for all $i \in [n]$.

The strength of this theorem obviously stands and falls with our ability to find suitable functions f_i. We will have a closer look at this issue in Sections 1.9 and 1.11. In Section 1.9 we give a sufficient condition for the existence of the f_i that gives rise to Corollaries 1.9.2 and 1.9.4. In Section 1.11 we suggest a concrete way of constructing the f_i.

1.7. The Topological Tool

The topological tool we are going to use in the proof of Theorem 1.6.1 is the Poincaré-Miranda Theorem. Let $\mathcal{S}^n$ denote the set of all partial functions $\sigma : [n] \to \{-1, 1\}$. Let $Q^n = [-1, 1]^n$ denote the n-dimensional cube. For any $\sigma \in \mathcal{S}^n$ let

$$Q^n_\sigma = Q^n \cap \{(x_1, \ldots, x_n) : x_i = \sigma(i) \ \forall i \in \operatorname{dom} \sigma\}$$

denote a (closed) face of Q^n. If $|\operatorname{dom} \sigma| = n$, Q^n_σ is a vertex. If $|\operatorname{dom} \sigma| = 1$, Q^n_σ is a facet. If $\operatorname{dom} \sigma = \emptyset$, $Q^n_\sigma = Q^n$.

1.7.1. Poincaré-Miranda Theorem.

Let $q : Q^n \to \mathbb{R}^n$ be a continuous function such that for all $\sigma \in \mathcal{S}^n$ and all $i \in \operatorname{dom} \sigma$

$$\sigma(i) q_i(Q^n_\sigma) \geq 0.$$

Then there exists an $x \in Q^n$ such that $q(x) = 0$.

Note that $\sigma(i) q_i(Q^n_\sigma) \geq 0$ if and only if the following holds: if $\sigma(i) = +1$ then $q_i(x) \geq 0$ for all $x \in Q^n_\sigma$ and if $\sigma(i) = -1$ then $q_i(x) \leq 0$ for all $x \in Q^n_\sigma$. Note also that we can, equivalently, restrict ourselves to those σ with $|\operatorname{dom} \sigma| = 1$ in the statement of the theorem.

The Poincaré-Miranda Theorem was shown in 1886 by Poincaré [Poi86]. In 1940 Miranda [Mir40] showed it to be equivalent to Brouwer's Fixpoint Theorem from 1911. It is a beautiful generalization of the Intermediate Value Theorem of Bolzano to higher dimensions, and it appears to be somewhat under-appreciated in modern topological geometry. For more information on this theorem we refer the reader to [Ist81, Chapter 4] and [Kul97].

The method by which we are going to apply this tool to the proof of our theorem is encapsulated in the following lemma which may be viewed as a slight generalization of the Poincaré-Miranda Theorem. It is stated and proved here in a more general version than is necessary for the proof of Theorem 1.6.1, in the hope that this method may be useful in other contexts as well.

1.7.2. Lemma.

Let X be a topological space, $p : X \to \mathbb{R}^n$ be a continuous function and $(C^*_\sigma)_{\sigma \in \mathcal{S}^n}$ a family of sets with the following properties:

1. $C^*_\sigma \subset X$ for all $\sigma \in \mathcal{S}^n$.
2. C^*_σ is $(n - |\operatorname{dom} \sigma| - 1)$-connected for all $\sigma \in \mathcal{S}^n$.
3. If $\sigma_1 \subset \sigma_2$, then $C^*_{\sigma_1} \supset C^*_{\sigma_2}$ for all $\sigma_1, \sigma_2 \in \mathcal{S}^n$.
4. $\sigma(i) p_i(C^*_\sigma) \geq 0$ for all $\sigma \in \mathcal{S}^n$ and all $i \in \operatorname{dom} \sigma$.

Then there is an element $y \in C^*_\emptyset$ such that $p(y) = 0$.

Proof. *Strategy.* We are going to show that there is a continuous map $q : Q^n \to C^*_\emptyset \subset X$ such that $\sigma(i)(p_i \circ q)(Q^n_\sigma) \geq 0$ for all $\sigma \in \mathcal{S}^n$ and all $i \in \mathsf{dom}\,\sigma$. Then we are done by the Poincaré-Miranda Theorem, as these conditions allow us to conclude that there exists an $x \in Q^n$ such that $(p \circ q)(x) = 0$, which means that $y := q(x) \in C^*_\emptyset$ is an element with $p(y) = 0$.

To show this claim we proceed by induction over $|\mathsf{dom}\,\sigma|$. For each $\sigma \in \mathcal{S}^n$ we are going to define $q|_{Q^n_\sigma}$ such that $q(Q^n_\sigma) \subset C^*_\sigma$. The foundation of the induction is the definition of $q(Q^n_\sigma)$ for σ with $\mathsf{dom}\,(\sigma) = [n]$, i.e. placing the corners of the cube.

Placing the corners. For each σ with $\mathsf{dom}\,\sigma = [n]$ the set Q^n_σ is a singleton $\{x_\sigma\}$. The set C^*_σ is -1-connected which means non-empty, so we can pick a point $y_\sigma \in C^*_\sigma$ arbitrarily and define $q(x_\sigma) := y_\sigma$. We now have a defined a (continuous) function q on the 0-skeleton of Q^n.

Induction over the size of $\mathsf{dom}\,\sigma$*.* Let σ be any sign function with $|\mathsf{dom}\,\sigma| = k$. Let $\sigma_1, \ldots, \sigma_l$ denote those sign functions with $\sigma_i \supset \sigma$ and $|\mathsf{dom}\,\sigma_i| = k + 1$. By induction $q|_{\bigcup_{i=1}^{l} Q^n_{\sigma_i}}$ is defined and continuous. We are now going to extend this definition to Q^n_σ.

Note that Q^n_σ is an $n - k$-dimensional ball and $\bigcup_{i=1}^{l} Q^n_{\sigma_i} = \partial Q^n_\sigma$ is an $n - k - 1$-dimensional sphere. Because $\sigma_i \supset \sigma$ we know by assumption that $C^*_{\sigma_i} \subset C^*_\sigma$. So $q(\bigcup_{i=1}^{l} Q^n_{\sigma_i}) \subset C^*_\sigma$. Because C^*_σ is $(n-k-1)$-connected, we can extend the definition of q continuously to Q^n_σ. We proceed in the same way with all σ with $|\mathsf{dom}\,\sigma| = k$ to obtain a continuous function q defined on the $n - k$-skeleton of Q^n.

We continue until we have defined q on ∂Q^n using all the σ with $|\mathsf{dom}\,\sigma| = 1$. The image of $q|_{\partial Q^n}$ lies in $C^*_\emptyset$ which is $n - 1$-connected, so we finish the construction by extending q continuously to all of Q^n.

Conclusion. q is continuous and defined on all of Q^n. Its image lies in $C^*_\emptyset \subset \mathsf{dom}\,p$. Now let $\sigma \in \mathcal{S}^n$ be any sign function. By assumption $\sigma(i)p_i(C^*_\sigma) \geq 0$ and as $q(Q^n_\sigma) \subset C^*_\sigma$ we conclude that $\sigma(i)(p_i \circ q)(Q^n_\sigma) \geq 0$ for all $i \in \mathsf{dom}\,\sigma$. So $p \circ q$ meets the conditions of the Poincaré-Miranda Theorem. □

1.8. Proof of Main Result

The strategy is to apply Lemma 1.7.2. To that end we will first define p and then go on to define sets C_σ as an intermediate step towards the definition of the sets C^*_σ.

Definition of p**.** A pair $(v, \lambda) \in \mathbb{R}^n \times \mathbb{R}$ defines an oriented hyperplane in $\mathbb{R}^n$ if $v \neq 0$. If $v = 0$ we have no interpretation for (v, λ). Hence we put $X = \mathbb{R}^n \setminus \{0\} \times \mathbb{R}$. Note that we can assume the f_i to be defined on all of $\mathbb{R}^n \setminus \{0\}$ and not only on S^{n-1} as we can extend them to the larger domain via $f_i(v) = f_i(v/|v|)$, keeping $\operatorname{Im} f_i$ unchanged. Now we define

$$\begin{aligned} p : X &\to \mathbb{R}^n \\ (v, \lambda) &\mapsto (\mu_i(H^+_{v,\lambda}) - \alpha_i)_{i \in [n]} \end{aligned}$$

which is a continuous function with the property that $p(v,\lambda) = 0$, if and only if $\mu_i(H^+_{v,\lambda}) = \alpha_i$ for all $i \in [n]$.

Definition and properties of C_σ. For each $\sigma \in \mathcal{S}^n$ with $|\operatorname{dom}\sigma| = n$ let H_σ denote a hyperplane with $\operatorname{int} H_\sigma^{\sigma(i)} \supset \operatorname{Im} f_i$ for all $i \in [n]$. Such a hyperplane exists by assumption. Let $\mathcal{H}'$ denote the arrangement of all these H_σ. The sets $\operatorname{Im} f_i$ are all contained in distinct n-dimensional cells of $\mathcal{H}'$. By assumption the $\operatorname{Im} f_i$ are bounded, so we can add hyperplanes to $\mathcal{H}'$ to obtain a hyperplane arrangement $\mathcal{H}$ in which each $\operatorname{Im} f_i$ is contained in a bounded n-dimensional cell. Let S_i denote the closed n-dimensional cell in $\mathcal{H}$ containing $\operatorname{Im} f_i$. The sets $S_1, \dots, S_n \subset \mathbb{R}^n$ are n-dimensional polytopes. Let $V(S_i)$ denote their vertex sets.

For all partial functions $\sigma : [n] \to \{-1, +1\}$ we define

$$C_\sigma := \{(v,\lambda) : \forall i \in \operatorname{dom}\sigma\ \forall x \in V(S_i) : \sigma(i)\langle v, x\rangle \geq \sigma(i)\lambda\}.$$

We now check whether properties 1. through 4. from Lemma 1.7.2 hold for the family $(C_\sigma)_{\sigma\in\mathcal{S}^n}$. First of all $(C_\sigma)_{\sigma\in\mathcal{S}^n}$ satisfies condition 4. as Lemma 1.8.1 shows.

1.8.1. Lemma.

If $(v,\lambda) \in C_\sigma$, then $\sigma(i)p_i(v,\lambda) \geq 0$ for all $i \in \operatorname{dom}\sigma$.

Proof. Let $(v,\lambda) \in C_\sigma$ and $i \in \operatorname{dom}\sigma$. For any $x \in \mathbb{R}^n$ the inequality $\sigma(i)\langle v,x\rangle \geq \sigma(i)\lambda$ holds, if and only if $x \in H^{\sigma(i)}_{v,\lambda}$. As the S_i are convex $H^{\sigma(i)}_{v,\lambda} \supset S_i \supset \operatorname{Im} f_i$ follows and hence $f_i(v) \in H^{\sigma(i)}_{v,\lambda}$ which means $H^{\sigma(i)}_{v,\lambda} \supset H_{v,\langle v,f_i(v)\rangle}$. If $\sigma(i) = +1$ this implies $H^+_{v,\lambda} \supset H^+_{v,\langle v,f_i(v)\rangle}$, so $\mu_i(H^+_{v,\lambda}) \geq \alpha_i$. If $\sigma(i) = -1$ this implies $H^+_{v,\lambda} \subset H^+_{v,\langle v,f_i(v)\rangle}$, so $\mu_i(H^+_{v,\lambda}) \leq \alpha_i$. □

That $(C_\sigma)_{\sigma\in\mathcal{S}^n}$ satisfies 3. is immediate as $\sigma_1 \subset \sigma_2$ means that any $(v,\lambda) \in C_{\sigma_1}$ has to satisfy a subset of the constraints defining C_{σ_2}. 2. also holds for the sets C_σ as it turns out that they are non-empty and convex and hence contractible. In fact we can show more:

1.8.2. Lemma.

C_σ is a polyhedral cone, for all $\sigma \in \mathcal{S}^n$. If $\operatorname{dom}\sigma \neq \emptyset$, then C_σ is pointed. For all $\sigma \in \mathcal{S}^n$ there exists $(v,\lambda) \in C_\sigma$ with $v \neq 0$.

Proof. *C_σ is a polyhedron.* C_σ is defined by a finite system of linear inequalities.

C_σ is a cone. For any $\beta \geq 0$ the implication

$$\sigma(i)\langle v,x\rangle \geq \sigma(i)\lambda \Rightarrow \sigma(i)\langle \beta v, x\rangle \geq \sigma(i)\beta\lambda$$

holds and hence if $(v,\lambda) \in C_\sigma$, then $\beta(v,\lambda) \in C_\sigma$.

C_σ is pointed. Suppose $i \in \mathsf{dom}\,\sigma$ and both (v,λ) and $-(v,\lambda)$ are in C_σ, then for all $x \in S_i$ we have

$$\sigma(i)\langle v,x\rangle \le \sigma(i)\lambda \quad \wedge \quad \sigma(i)\langle v,x\rangle \ge \sigma(i)\lambda$$

so in fact equality holds. This implies that S_i lies in the hyperplane defined by this equation, which is a contradiction to S_i being n-dimensional.

There exists $(v,\lambda) \in C_\sigma$ with $v \neq 0$. Let $\sigma \in \mathcal{S}^n$. Let S_i' denote the n-dimensional cells of $\mathcal{H}'$ containing Im f_i for all $i \in [n]$. By definition of $\mathcal{H}'$ there is a hyperplane H such that $H^{\sigma(i)} \supset S_i' \supset S_i$ for all $i \in \mathsf{dom}\,\sigma$. □

Note that $(0,0) \in C_\sigma$ for all $\sigma \in \mathcal{S}^n$, $\{(0,\lambda) : \lambda \ge 0\} \subset C_\sigma$ iff Im $\sigma \subset \{-1\}$ and $\{(0,\lambda) : \lambda \le 0\} \subset C_\sigma$ iff Im $\sigma \subset \{+1\}$. So the sets C_σ meet conditions 2., 3. and 4., but they do not meet condition 1. However, if we were to correct this by removing $\{(0,\lambda) : \lambda \in \mathbb{R}\}$ we would destroy the connectivity of the C_σ and violate condition 2. Note that X itself is only $n-2$-connected and not $n-1$-connected. So we have to use a different construction.

Definition and properties of the C_σ^*. For each $i \in [n]$, let b_i be the barycenter of S_i. Let H_{v_0,λ_0} be an oriented hyperplane containing all the b_i. This exists, because $b_1,\dots,b_n$ are only n points in $\mathbb{R}^n$. Define

$$C_\sigma^* := C_\sigma \setminus \{(\beta v_0,\lambda) : \lambda \in \mathbb{R}, \beta \ge 0\}.$$

Note that $\{(\beta v_0,\lambda) : \lambda \in \mathbb{R}, \beta \ge 0\} \supset \{(0,\lambda) : \lambda \in \mathbb{R}\}$ and hence $C_\sigma^* \subset X$, so these sets fulfill condition 1. As $C_\sigma^* \subset C_\sigma$, condition 4. still holds. As we removed the same set from all the C_σ, we also have condition 3. What is not immediately clear, is whether 2. holds. We are going to show that the C_σ^* are contractible, and hence $(n - |\mathsf{dom}\,\sigma| - 1)$-connected as required.

1.8.3. Lemma.

For any $\sigma \in \mathcal{S}^n$ the set C_σ^* is contractible.

Proof. *Case 1:* $|\text{Im}\ \sigma| = 2$. We are going to show that C_σ^* is non-empty and convex, which implies that C_σ^* is contractible.

We start out by showing that $C_\sigma^* = C_\sigma \cap X$. Suppose that $(\beta v_0,\lambda) \in C_\sigma$ for some $\beta > 0$. Let, say, $\sigma(1) = +1$ and $\sigma(2) = -1$. Then $b_1 \in S_1 \subset H^+_{\beta v_0,\lambda}$ and $b_2 \in S_2 \subset H^-_{\beta v_0,\lambda}$, so $\lambda \le \langle \beta v_0, b_1\rangle = \langle \beta v_0, b_2\rangle \le \lambda$ where the equality holds by the choice of v_0. So $b_1, b_2 \in H_{\beta v_0,\lambda}$. But the S_i are n-dimensional. So there exist $x_1, x_2 \in S_1$ with $\langle \beta v_0, x_1\rangle < \lambda < \langle \beta v_0, x_2\rangle$, which is a contradiction to $S_1 \subset H^+_{\beta v_0,\lambda}$.

We now know that $C_\sigma^* = C_\sigma \cap X$ and as we have already seen, C_σ contains a pair (v,λ) such that $v \neq 0$, so we can conclude that $(v,\lambda) \in C_\sigma^*$. Also we know that C_σ is convex and we want to show that C_σ^* is convex as well. It suffices to argue that there are no v, λ, λ' such that both (v,λ) and $(-v,\lambda')$ are in C_σ. Assume to the contrary that this is the case. Then

$$H^+_{v,\lambda} \supset S_1 \subset H^+_{-v,\lambda'} \text{ and } H^-_{v,\lambda} \supset S_2 \subset H^-_{-v,\lambda'}.$$

But

$$H^+_{-v,\lambda'} = H^-_{v,-\lambda'} \text{ and } H^-_{-v,\lambda'} = H^+_{v,-\lambda'},$$

so

$$S_1 \subset H^+_{v,\lambda} \cap H^-_{v,-\lambda'} \text{ and } S_2 \subset H^-_{v,\lambda} \cap H^+_{v,-\lambda'},$$

and hence

$$\forall s_1 \in S_1 : \lambda \le \langle v, s_1 \rangle \le -\lambda' \text{ and } \forall s_2 \in S_2 : -\lambda' \le \langle v, s_2 \rangle \le \lambda.$$

As both S_1 and S_2 are non-empty, we conclude that $\lambda = -\lambda'$ which implies

$$S_1 \subset H^+_{v,\lambda} \cap H^-_{v,\lambda} = H_{v,\lambda} \text{ and } S_2 \subset H^-_{v,\lambda} \cap H^+_{v,\lambda} = H_{v,\lambda}.$$

But the S_i are full dimensional and cannot be contained in a hyperplane, which is a contradiction.

Case 2: $|\text{Im } \sigma| \le 1$. Without loss of generality, we assume that Im $\sigma \subset \{+1\}$. σ may be empty. For $-v_0$ there exists a λ' such that $H^+_{-v_0,\lambda'} \supset S_i$ for all i, because the S_i are bounded. This shows that C^*_σ is non-empty and $(-v_0, \lambda') \in C^*_\sigma$. We are going to give a deformation retraction of C^*_σ to $(-v_0, \lambda')$. This means that we are looking for a homotopy $h_t : C^*_\sigma \to C^*_\sigma$ such that $h_0 = \text{id}$ and h_1 is the map with $h_1(v, \lambda) = (-v_0, \lambda')$ for all (v, λ). We are going to use the straight line homotopy, defined by

$$h_t(v, \lambda) = t(-v_0, \lambda') + (1 - t)(v, \lambda).$$

h_t is continuous, $h_0(v, \lambda) = (v, \lambda)$ and $h_1(v, \lambda) = (-v_0, \lambda')$ as desired. We still have to show that Im $h_t \subset C^*_\sigma$ for all t. C_σ is a cone and hence convex, so all we have to argue is that no $(\beta v_0, \gamma)$ with $\beta \ge 0$ and $\gamma \in \mathbb{R}$ lies in Im h_t for any t. We have already seen that Im $h_0 = C^*_\sigma$ and Im $h_1 = \{(-v_0, \lambda')\}$, so we only have to consider $0 < t < 1$. Suppose $h_t(v, \lambda) = (\beta v_0, \gamma)$ for such β and γ for some $t \in (0, 1)$. Then

$$-tv_0 + (1 - t)v = \beta v_0$$

which we can solve for v to obtain

$$v = \frac{\beta + t}{1 - t} v_0.$$

We have thus written v as a non-negative multiple of v_0. But then $(v, \lambda) \notin C^*_\sigma$. So C^*_σ is contractible. □

We have now seen that the C^*_σ also meet condition 2. and hence all the requirements of Lemma 1.7.2. Applying Lemma 1.7.2 gives us the desired $(v, \lambda) \in \mathbb{R}^{n+1}$ with $v \ne 0$ such that $p(v, \lambda) = 0$. This concludes the proof of Theorem 1.6.1.

1.9. *Existence of the f_i*

1.9.1. Proposition.

Let μ be a continuous probability measure and $\alpha \in [0,1]$. If S is a compact connected set with $\mu(S) \geq \max\{\alpha, 1-\alpha\}$, then there exists a function $f : S^{n-1} \to \mathbb{R}^n$ such that $\mu(H^+_{v,\langle v,f(v)\rangle}) = \alpha$ for all $v \in S^{n-1}$ and $\operatorname{Im} f \subset S$.

Proof. First of all, note that we can assume $\alpha \leq \frac{1}{2}$ without loss of generality. By definition a set S with the given properties exists for a parameter α if and only if it exists for a parameter $1-\alpha$. Given an auxiliary function f with $\mu(H^+_{v,\langle v,f(v)\rangle}) = \alpha$ for all $v \in S^{n-1}$ and $\operatorname{Im} f \subset S$, we can obtain an auxiliary function f' with $\mu(H^+_{v,\langle v,f'(v)\rangle}) = 1-\alpha$ for all $v \in S^{n-1}$ and $\operatorname{Im} f' \subset S$ by putting $f'(v) = f(-v)$. Therefore, let $\alpha \leq \frac{1}{2} \leq 1-\alpha$.

Let $v \in S^{n-1}$. If we can show that there exists a hyperplane $H_{v,\lambda}$ with $\mu(H^+_{v,\lambda}) = \alpha$ and $H_{v,\lambda} \cap S \neq \emptyset$ we are done. S is compact and hence bounded. So there exist hyperplanes H_{v,λ_+} and H_{v,λ_-} such that $\mu(H^+_{v,\lambda_+}) \geq 1-\alpha \geq \alpha \geq \mu(H^+_{v,\lambda_-})$. By continuity of μ there exists a hyperplane H_{v,λ_0} with $\mu(H^+_{v,\lambda_0}) = \alpha$. Now we distinguish three cases.

Case 1: $H^+_{v,\lambda_0} \supset S$. Then $\alpha = \mu(H^+_{v,\lambda_0}) \geq \mu(S) \geq 1-\alpha$ and hence $\alpha = 1-\alpha$ as well as $\mu(H^+_{v,\lambda_0}) = \mu(S)$. Put $\lambda = \sup\{\lambda' \geq \lambda_0 : H^+_{v,\lambda'} \supset S\}$. As S is compact $H_{v,\lambda} \cap S \neq \emptyset$.

$$\alpha = \mu(H^+_{v,\lambda_0}) \geq \mu(H^+_{v,\lambda'}) \geq \mu(S) \geq 1-\alpha = \alpha$$

for every $\lambda' \geq \lambda_0$ with $H^+_{v,\lambda'} \supset S$, so by continuity of μ we also have $\mu(H^+_{v,\lambda}) = \alpha$.

Case 2: $H^-_{v,\lambda_0} \supset S$. Then $1-\alpha = \mu(H^-_{v,\lambda_0}) \geq \mu(S) \geq 1-\alpha$ and hence $\mu(H^-_{v,\lambda_0}) = \mu(S)$. Put $\lambda = \inf\{\lambda' \leq \lambda_0 : H^-_{v,\lambda'} \supset S\}$. As S is compact $H_{v,\lambda} \cap S \neq \emptyset$.

$$1-\alpha = \mu(H^-_{v,\lambda_0}) \geq \mu(H^-_{v,\lambda'}) \geq \mu(S) \geq 1-\alpha$$

for every $\lambda' \geq \lambda_0$ with $H^-_{v,\lambda'} \supset S$, so by continuity of μ we have $\mu(H^+_{v,\lambda}) = 1-\mu(H^-_{v,\lambda}) = \alpha$.

Case 3: Neither inclusion holds. Then $H^+_{v,\lambda_0} \cap S \neq \emptyset \neq H^-_{v,\lambda_0} \cap S$ and because S is connected $H_{v,\lambda_0} \cap S \neq \emptyset$. □

If $0 \neq \alpha \neq 1$ or the support of μ is bounded, then a set S as in Proposition 1.9.1 always exists. There are examples where the support of μ is unbounded such that for $\alpha \in \{0,1\}$ no function f with $\mu(H^+_{v,\langle v,f(v)\rangle}) = \alpha$ exists.

1.9.2. Corollary.

Let $\mu_1, \ldots, \mu_n$ be continuous probability measures on $\mathbb{R}^n$. Let $\alpha_1, \ldots, \alpha_n \in [0,1]$. Let $S_1, \ldots, S_n \subset \mathbb{R}^n$ have one of the following two properties:

1. $S_1, \ldots, S_n$ are compact, can be separated and
$$\mu_i(S_i) \geq \max\{\alpha_i, 1-\alpha_i\} \text{ for all } i \in [n].$$
2. $S_1, \ldots, S_n$ are closed, can be separated and
$$\mu_i(S_i) > \max\{\alpha_i, 1-\alpha_i\} \text{ for all } i \in [n].$$

Then there exists a hyperplane H such that $\mu_i(H^+) = \alpha_i$ for all $i \in [n]$.

Proof. Suppose $S_1, \ldots, S_n$ are compact, can be separated and $\mu_i(S_i) \geq \max\{\alpha_i, 1-\alpha_i\}$ for all $i \in [n]$. We first pass to the respective convex hulls $S_i' = \mathsf{conv}(S_i)$ which are compact, convex and hence connected, can be separated and have $\mu_i(S_i') \geq \max\{\alpha_i, 1-\alpha_i\}$. Applying 1.9.1 and 1.6.1 shows that condition 1. suffices for the existence of the desired hyperplane.

If $S_1, \ldots, S_n$ are closed, can be separated and $\mu_i(S_i) > \max\{\alpha_i, 1-\alpha_i\}$ for all i, then for every i there exists a ball B_i such that $S_i' := S_i \cap B_i$ is compact and $\mu_i(S_i') > \max\{\alpha_i, 1-\alpha_i\}$. $S_1', \ldots, S_n'$ can be separated. Using the sufficiency of 1. we obtain that 2. is sufficient for the existence of the desired hyperplane. ☐

As an application we present this very simple consequence about two measures.

1.9.3. Corollary.

Let μ_1, μ_2 be two probability measures on $\mathbb{R}^2$ and $\alpha_1, \alpha_2 \in (0,1)$. If there exists a hyperplane H such that
$$\mu_1(H^+) > \max\{\alpha_1, 1-\alpha_1\} \quad \text{and} \quad \mu_2(H^-) > \max\{\alpha_2, 1-\alpha_2\},$$
then there exists an (α_1, α_2)-splitting of (μ_1, μ_2).

Proof. As $\mu_1(H^+) > \max\{\alpha_1, 1-\alpha_1\}$ and $0 < \alpha_1 < 1$ there is a compact subset $S_1 \subset \mathsf{int}\, H^+$ such that $\mu_1(S_1) > \max\{\alpha_1, 1-\alpha_1\}$. Similarly there is a compact subset $S_2 \subset \mathsf{int}\, H^-$ such that $\mu_2(S_2) > \max\{\alpha_2, 1-\alpha_2\}$. Let H' be a hyperplane such that the open positive half-space determined by H' contains both S_1, S_2. Then the hyperplanes H and H' together with the respective hyperplanes with opposite orientation show that S_1 and S_2 can be separated. Thus by Corollary 1.9.2 there exists (α_1, α_2)-splitting of (μ_1, μ_2). ☐

Corollary 1.9.2 can be simplified further. Let S denote the closure of the support of μ. If S is bounded, then S fulfills the conditions of Corollary 1.9.2 for any choice of α. This immediately gives rise to the following result.

1.9.4. Corollary.

Let $\mu_1, \ldots, \mu_n$ be continuous probability measures on $\mathbb{R}^n$. Let $S_1, \ldots, S_n$ denote the closure of their respective support.

If $S_1, \ldots, S_n$ are bounded and can be separated, then for any ratios $\alpha_1, \ldots, \alpha_n \in [0,1]$, there exists a hyperplane H such that $\mu_i(H^+) = \alpha_i$ for all $i \in [n]$.

As mentioned in the introduction of this chapter, Bárány, Hubard and Jerónimo gave - recently and independently - a proof of this result in the special case that the measures μ_i have the following additional property: Let μ be any one of the μ_i. For any $v \in \mathbb{R}^n \setminus \{0\}$ let $\lambda_- = \inf\{\lambda : \mu(H^+_{v,\lambda}) > 0\}$ and $\lambda_+ = \sup\{\lambda : \mu(H^+_{v,\lambda}) < 1\}$. Then

$$\mu(\{x \in \mathbb{R}^n : \lambda_1 \leq \langle x, v \rangle \leq \lambda_2\}) > 0$$

for any $\lambda_1, \lambda_2 \in [\lambda_-, \lambda_+]$ with $\lambda_1 < \lambda_2$. Corollary 1.9.4 does not impose this requirement on the μ_i. The proof of Bárány et al. is different from our proof. It relies on Brouwer's Fixpoint Theorem instead of the Poincaré-Miranda Theorem.

Also as mentioned in the introduction, Stojmenović remarks in [Sto90] that his methods can be generalized to obtain a statement similar to Corollary 1.9.4. However, no proof is given and in the journal version [Sto91] of [Sto90], that remark is weakened to include only the case $n = 3$. Again the method employed by Stojmenović is different from ours.

It should be noted that Corollary 1.9.2 is strictly more general than Corollary 1.9.4, as Corollary 1.9.2 does not require the supports of the μ_i to be disjoint. Also Theorem 1.6.1 is strictly more general than Corollary 1.9.2, as the convex hulls of the sets Im f_i can often be chosen to be strict subsets of the sets S_i. Thus 1.6.1 and 1.9.2 give partial answers to a question in [Sto91] who specifically asked about methods that can be employed if the supports of the μ_i overlap.

1.10. Partitioning one Mass by two Hyperplanes

It is an easy and well-known corollary of the Ham Sandwich Theorem in the plane, that any continuous probability measure in the plane can be partitioned by two lines into four parts of measure $\frac{1}{4}$. See e.g. [GH05], [Meg85], [Mat08]. In just the same way Corollary 1.9.2 implies the following theorem, which the author was unable to find in the prior literature.

1.10.1. Theorem.

For any continuous probability measure μ on $\mathbb{R}^n$ and any $\alpha_1, \ldots, \alpha_4 \in (0,1)$ with $\sum_i \alpha_i = 1$, there exist two oriented hyperplanes H_1, H_2 such that $H_1^+ \cap H_2^+$, $H_1^- \cap H_2^+$, $H_1^+ \cap H_2^-$, $H_1^- \cap H_2^-$ have μ-measure $\alpha_1, \ldots, \alpha_4$, respectively.

Proof. Without loss of generality, we can assume $n = 2$. For higher n we pick an arbitrary projection down to $\mathbb{R}^2$. Now we pick any normal v and let $H_{v,\lambda} =: H_2$ denote a hyperplane with $\mu(H^+_{v,\lambda}) = \alpha_1 + \alpha_2$. Now define

$$\mu_1(A) := \frac{1}{\alpha_1 + \alpha_2}\mu(A \cap H^+_{v,\lambda}) \quad \text{and} \quad \mu_2(A) := \frac{1}{\alpha_3 + \alpha_4}\mu(A \cap H^-_{v,\lambda}).$$

μ_1 and μ_2 are again continuous probability measures. Because all α_i are non-zero, there exist $\lambda_1 > \lambda > \lambda_2$ such that

$$\mu_1(H^+_{v,\lambda_1}) > \max\{\frac{\alpha_1}{\alpha_1 + \alpha_2}, 1 - \frac{\alpha_1}{\alpha_1 + \alpha_2}\}$$

and

$$\mu_2(H^-_{v,\lambda_2}) > \max\{\frac{\alpha_3}{\alpha_3 + \alpha_4}, 1 - \frac{\alpha_3}{\alpha_3 + \alpha_4}\}.$$

H^+_{v,λ_1} and H^-_{v,λ_2} are closed and they are separated by $H_{v,\lambda}$. Applying Corollary 1.9.2 yields a hyperplane H_1 such that $\mu_1(H^+_1) = \frac{\alpha_1}{\alpha_1+\alpha_2}$ and $\mu_2(H^+_1) = \frac{\alpha_3}{\alpha_3+\alpha_4}$. By construction these two hyperplanes H_1, H_2 now have the desired properties. $\square$

Note that we can even prescribe the normal of one of the two hyperplanes. We cannot, though, control the angle in which the hyperplanes meet. It is a result by Zindler [Zin21] that any convex object can be partitioned into four parts of equal area by two lines that are *perpendicular* to each other. A related problem was posed by Birch (see [GH05]): For any convex object and any $0 < \alpha < \frac{1}{2}$ are there two *perpendicular* lines that partition the object into four pieces having a fraction $\alpha, \alpha, 1 - \alpha, 1 - \alpha$, respectively, of the total area? This problem is still open.

1.11. Central Spheres

The method given in Section 1.9 allows us to control Im f_i, but it gives a lot of freedom for choosing the particular function f_i. What is a good canonical choice? We suggest what we call the central sphere, which we define below. First we need to recall some definitions. Given a density function h on $\mathbb{R}^k$, the **center of mass** of h is defined as

$$\frac{1}{M}\int_{\mathbb{R}^k} \mathrm{id} \cdot h \, dL$$

where $M = \int_{\mathbb{R}^k} h \, dL$ and L is the Lebesgue measure on $\mathbb{R}^k$. Given a set $S \subset \mathbb{R}^k$, we consider the affine subspace A of minimal dimension that contains S. Let L' denote the Lebesgue measure on A. The **centroid** of S is defined as

$$\frac{1}{L'(S)}\int_S \mathrm{id} \, dL'$$

which we interpret as a point in $\mathbb{R}^k$.

Let μ denote a continuous probability measure and h a fixed density function of μ (see Section 1.2). Let $\alpha \in [0,1]$. Let S be a convex set that is compact and has $\mu(S) \geq \max\{\alpha, 1-\alpha\}$. We now define an auxiliary function $c : S^{n-1} \to \mathbb{R}^n$ which we call the **central sphere** of μ, h, S and α as follows. For each $v \in S^{n-1}$ there exists a hyperplane $H_{v,\lambda}$ with $\mu(H^+_{v,\lambda}) = \alpha$ and $H_{v,\lambda} \cap S \neq \emptyset$ as described in the proof of Proposition 1.9.1. Let χ_S denote the characteristic function of S. If $(h \cdot \chi_S)|_{H_{v,\lambda}} = 0$ almost everywhere with respect to the Lebesgue measure on $H_{v,\lambda}$, we define $c(v)$ to be the centroid of $H_{v,\lambda} \cap S$. Otherwise we define $c(v)$ to be the center of mass of $(h \cdot \chi_S)|_{H_{v,\lambda}}$ with respect to the affine space $H_{v,\lambda}$. With this definition $\mu(H^+_{v,\langle v,c(v)\rangle}) = \alpha$ for all $v \in S^{n-1}$ and $\operatorname{Im} c \subset S$.

This is just the construction we used in Example 1.3.1: If B is a ball, $\mu(\Omega) = \frac{L(\Omega)}{L(B)}$, $h = \chi_B$, $S = B$ and $0 < \alpha < \frac{1}{2}$, then the auxiliary function f defined in Example 1.3.1 is the central sphere of μ, h, S, α. Note that central spheres are also related to a construction used by Bárány et al. in [BHJ08].

The central spheres are particular functions that can be used in the proof of Corollary 1.9.2. They give rise to a slightly stronger version of that corollary, in which we only need to require that the sets $\operatorname{Im} c_i$ can be separated, not the sets S_i. (Here c_i denotes the central sphere of μ_i, h_i, S_i and α_i.) Note that this gains us something only if we fix a density function h_i for each μ_i a priori: for each $\mu_1, \ldots, \mu_n$ and $S_1, \ldots, S_n$ there is a choice of density functions $h_1, \ldots, h_n$ such that if $\operatorname{Im} c_1, \ldots, \operatorname{Im} c_n$ can be separated, then $S_1, \ldots, S_n$ can be separated.

1.11.1. Example. Consider a regular 5-gon P in the plane. Let $\mu(\Omega) = \frac{L(\Omega)}{L(P)}$, $h = \chi_P$ and $S = P$. We consider the central spheres for $\alpha = \frac{1}{2}$ and for some small $\alpha > 0$. Qualitative pictures of the corresponding central spheres are shown in Figure 1.4, see also [DEBG$^+$06]. There are a couple of things to note. 1) In the case $\alpha = \frac{1}{2}$ the central sphere has turning number 4 (not 2), while in the case of small α the curve has turning number 1. 2) Even in the case $\alpha = \frac{1}{2}$ the hyperplanes $H_{v,\lambda}$ with $\mu(H^+_{v,\lambda}) = \alpha$ do not in general contain the centroid of P. So even in this case, the f_i cannot be chosen to be constant. See [BCJ01]. 3) The image of the central sphere is not in general the boundary of a convex set. Notice how the "arcs" of the central sphere of the 5-gon shown in Figure 1.4 are bent inwards slightly.

Some remarks about the continuity of c: The central sphere of μ, α is not in general continuous, not even if the density h of μ is smooth. Moreover, even if we only define c for those v where $h|_H = 0$ does not hold almost a.e., c cannot be extended continuously to all v. An example for this is given in Figure 1.5: On each of the three disks, the density of μ is an identical smooth cap. Now, if $\alpha = \frac{1}{3}$ and the normal v points directly upwards, then $H := H_{v,\lambda(v)}$ is as shown in the figure. On H the density h is identical 0. Tilting v slightly to the left by *any* small amount, we always obtain a hyperplane that intersects H in p. Tilting v slightly to the right by *any* small amount, we always obtain a hyperplane that intersects H in q. No matter how we choose $c|_H(v)$, we can never obtain a continuous function.

So central spheres are an interesting object whose study is worthwhile. However, restricting the choice of f_i in Theorem 1.6.1 to central spheres, or to central spheres of the sets S_i in

Figure 1.4: Central spheres of a 5-gon for $\alpha = \frac{1}{2}$ (with 5-gon and magnified) and small $\alpha > 0$.

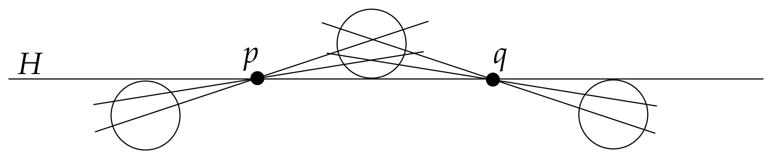

Figure 1.5: Central spheres are not continuous in general.

Corollary 1.9.2, significantly reduces its generality. Thus, in the context of Theorem 1.6.1, central spheres should be seen as a mainly illustrative device.

Chapter 2.

Staircases in $\mathbb{Z}^2$

Motivated by the study of lattice points inside polytopes, in this chapter we seek to understand the set of lattice points "close" to a rational line in the plane. To this end we define a staircase in the plane to be the set of lattice points in the half-plane below a rational line that have Manhattan Distance less than 1 to the line. We prove several properties of these point sets, most importantly we show that they have a recursive structure that is reminiscent of the Euclidean Algorithm, which leads to recursive descriptions of staircases, the lattice points in fundamental parallelepipeds of planar cones and the lattice points in certain planar triangles.

Not surprisingly, staircases are closely related to the Beatty and Sturmian sequences defined in number theory (see [Sto76, FMT78, PS90, O'B02]), i.e. to sequences of the form $\left(\left\lfloor \frac{b}{a}n \right\rfloor - \left\lfloor \frac{b}{a}(n-1) \right\rfloor\right)_n$ for $a, b \in \mathbb{N}$ with $\gcd(a,b) = 1$. We show several elementary properties of these sequences from a geometric point of view. To our knowledge such a geometric approach to these sequences is not available in the prior literature. Our observations lead to three characterizations of these sequences (Theorem 2.4.1). One of these is known (see [GLL78, Fra05]) while the other two are new.

We conclude the chapter by giving several applications of our findings. Firstly, we give a new proof of a theorem by Barvinok in dimension 2. Barvinok's Theorem states that the generating function of the lattice points inside a rational simplicial cone can be written as a short rational function. While Barvinok uses a signed decomposition of the cone into unimodular cones to achieve this result, we partition the cone into sets that have a short representation.

Secondly, these ideas can also be used to give a recursion formula for Dedekind-Carlitz polynomials. These are polynomials of the form $\sum_{k=1}^{a-1} x^{k-1} y^{\left\lfloor \frac{b}{a}k \right\rfloor}$ or, equivalently, generating functions of the lattice points inside the open fundamental parallelepipeds of cones in the plane. Our recursion formula answers a question from [BHM08].

Finally, we simplify ideas from [Sca85] and [Rez06] to give a partially new proof of White's Theorem, which characterizes three-dimensional lattice simplices that contain no lattice points except their vertices.

This chapter is organized as follows. In Section 2.1 we recall some preliminary definitions from linear algebra and the theory of integer points in polyhedra. In Section 2.2 we introduce staircases and Beatty and Sturmian sequences and present elementary facts about them. While Section 2.2 is rather concise, we elaborate more in Section 2.3, where staircases are examined from a geometric point of view. Most importantly we explain the recursive structure of staircases in Theorems 2.3.9, 2.3.11 and 2.3.12. In Section 2.4 we motivate the three characterizations of Sturmian sequences before summarizing them in Theorem 2.4.1. Section 2.5 is devoted to the proof of this theorem. In Section 2.6 we apply Theorem 2.3.12 to give a new proof of Barvinok's Theorem in dimension 2 and in Section 2.7 we use Theorem 2.3.11 to give a recursion formula for Dedekind-Carlitz sums. We conclude the chapter in Section 2.8 by giving a partially new proof of White's Theorem.

The results in this chapter are joint work with Frederik von Heymann.

2.1. *Linear Algebra and Lattice Points*

Before we introduce staircases, we give some preliminary definitions.

For any real number $r \in \mathbb{R}$ we define the **integral part** $\lfloor r \rfloor := \max\{z \in \mathbb{Z} \mid z \leq r\}$ of r. The **fractional part** $\{r\}$ of r is then defined by $r = \lfloor r \rfloor + \{r\}$. Given $0 < a, b \in \mathbb{N}$ there exist unique integers $(b \text{ div } a)$ and $(b \text{ mod } a)$ such that $b = (b \text{ div } a) \cdot a + (b \text{ mod } a)$ and $0 \leq b \text{ mod } a < a$. Using these two functions we can write $\left\lfloor \frac{b}{a} \right\rfloor = b \text{ div } a$ and $\{\frac{b}{a}\} = \frac{b \text{ mod } a}{a}$. We are going to use these two notations interchangeably.

Given $A, B \subset \mathbb{R}^2$ and $v \in \mathbb{R}^2$, we define $A + B := \{a + b \mid a \in A, b \in B\}$ and $-A := \{-a \mid a \in A\}$ and we use the abbreviations $A - B := A + (-B)$ and $A + v = A + \{v\}$. We will refer to $A + B$ as the **Minkowski sum** of A and B. The difference of sets is denoted with $A \setminus B := \{a \in A \mid a \notin B\}$. The standard basis vectors in $\mathbb{R}^n$ are denoted by $e_1, \dots, e_n$. We will sometimes index the dimensions of our vector spaces (or more generally the components of a product) with elements from a set S and write $\mathbb{R}^S$. In this case the standard basis vectors are denoted by e_s for $s \in S$. Moreover we identify X^Y with the set of all functions from Y into X. Thus an element $z \in \mathbb{R}^S$ is both a vector and a function $z : S \to \mathbb{R}$ and we will use both the notations z_s and $z(s)$ to describe the value of z at (entry) s. The **support** of z is the set $\mathsf{supp}(z) = \{s \in S \mid z(s) \neq 0\}$ and the **zero set** of z is its complement $\mathsf{zero}(z) = \{s \in S \mid z(s) = 0\}$.

We denote the all-zero vector by 0 and the all-one vector by 1. On $\mathbb{R}^S$, the relation $\leq$ denotes the **dominance order**, i.e. $v \leq w$ for vectors $v, w \in \mathbb{R}^S$ if and only if $v_s \leq w_s$ for all $s \in S$. Thus for a vector $v \in \mathbb{R}^n$ the statement $0 \leq v \leq 1$ is equivalent to saying that $v \in [0, 1]^n$ is an element of the n-dimensional unit cube.

Given a points $v_1, \dots, v_n, z \in \mathbb{R}^d$ we call a vector $(\lambda_1, \dots, \lambda_n) \in \mathbb{R}^n$ a **representation** of z **in terms of** the v_i if $\sum_{i=1}^{n} \lambda_i v_i = z$. The representation is **integral** if all of the λ_i are integers.

The representation is **non-negative** (**positive**, etc.) if all of the λ_i are non-negative (positive, etc.), respectively. The representation is **affine** if $\sum_{i=1}^{n} \lambda_i = 1$ and it is **convex** if it is affine and non-negative. Alternatively we say that z is an (integral, affine, etc.) **combination** of the v_i.

The integer points in $\mathbb{R}^d$ are called **lattice points** and the set $\mathbb{Z}^d$ of all these is called the **lattice**. A **lattice polytope** is a polytope whose vertices are lattice points. A **lattice basis** is a set B of lattice points such that every lattice point can be represented as an integral combination of the points in B. An affine (linear) **lattice transformation** of the plane is an affine (linear) automorphism of the plane that maps $\mathbb{Z}^2$ bijectively onto $\mathbb{Z}^2$. Two polytopes P_1, P_2 are **lattice equivalent** or **lattice isomorphic** if there exists an affine lattice transformation h such that $h(P_1) = P_2$, in which case we write $P_1 \approx P_2$. Whenever we say that a polytope $P \in \mathbb{R}^n$ and a polytope $Q \in \mathbb{R}^m$ for $m \geq n$ are lattice equivalent, we mean that P and Q are lattice equivalent if P is viewed as a polytope in $\mathbb{R}^m$ via the canonical inclusion $\mathbb{R}^n \to \mathbb{R}^m$. A vector $v \in \mathbb{Z}^d$ is **primitive** if $\gcd(v_1, \ldots, v_d) = 1$ or, equivalently, if $\mathbb{Z}^d \cap \mathsf{conv}(0, v) = \{0, v\}$.

For any set S in $\mathbb{R}^n$ the **Ehrhart** function L_S of S counts the number of lattice points in dilates of S, that is $\mathsf{L}_s(k) = \mathbb{Z}^n \cap kS$ for all $k \in \mathbb{Z}_{\geq 0}$. Note that $\mathsf{L}_s(0) = 1$. By Ehrhart's Theorem [Ehr77], if S is a lattice polytope then $\mathsf{L}_S(k)$ is a polynomial in k. For an introduction to these concept we recommend the book [BR07].

The d-dimensional **standard simplex** in $\mathbb{R}^d$ is

$$\mathsf{conv}\{0, e_1, \ldots, e_d\} = \{x \in \mathbb{R}^d : 0 \leq x, \sum_{i=1}^{d+1} x_i \leq 1\}.$$

The d-dimensional **standard simplex** in $\mathbb{R}^{d+1}$ is

$$\mathsf{conv}\{e_1, \ldots, e_{d+1}\} = \{x \in \mathbb{R}^{d+1} : 0 \leq x, \sum_{i=1}^{d+1} x_i = 1\}.$$

Which of the two we mean will be clear from context in any given instance. Note that when viewing both as d-simplices in $\mathbb{R}^{d+1}$, the two are lattice equivalent.

A n-dimensional simplex is **unimodular** if it is lattice equivalent to the n-dimensional standard simplex. An n-dimensional lattice simplex in $\mathbb{R}^n$ with vertices $v_0, \ldots, v_n \in \mathbb{R}^n$ is unimodular if and only if the matrix with columns $v_1 - v_0, \ldots, v_n - v_0$ has determinant ± 1. We need several other characterizations of unimodular lattice simplices. First of all a lattice simplex σ in $\mathbb{R}^d$ is unimodular if and only if every lattice point in its affine hull has an integral affine representation in terms of its vertices, i.e. for every $z \in \mathbb{Z}^d \cap \mathsf{aff}(\sigma)$ and every $\lambda \in \mathbb{R}^{\mathsf{vert}(\sigma)}$ with $z = \sum_{v \in \mathsf{vert}(\sigma)} \lambda_v v$ and $\sum_{v \in \mathsf{vert}(\sigma)} \lambda_v = 1$, we have $\lambda \in \mathbb{Z}^{\mathsf{vert}(\sigma)}$. This characterization can be varied in several ways: A lattice simplex σ in $\mathbb{R}^d$ is unimodular if and only if for every $z \in \mathbb{Z}^d$, every $k \in \mathbb{N}$ and every $\lambda \in \mathbb{R}^{\mathsf{vert}(\sigma)}$ with $z = \sum_{v \in \mathsf{vert}(\sigma)} \lambda_v v$ and $\sum_{v \in \mathsf{vert}(\sigma)} \lambda_v = k$, we have $\lambda \in \mathbb{Z}^{\mathsf{vert}(\sigma)}$. Moreover, a lattice simplex σ in $\mathbb{R}^d$ is unimodular if and only if for every $z \in \mathbb{Z}^d$, every $k \in \mathbb{N}$ and every $\lambda \in \mathbb{R}_{\geq 0}^{\mathsf{vert}(\sigma)}$ with $z = \sum_{v \in \mathsf{vert}(\sigma)} \lambda_v v$ and $\sum_{v \in \mathsf{vert}(\sigma)} \lambda_v = k$, we have $\lambda \in \mathbb{Z}^{\mathsf{vert}(\sigma)}$.

2.2. *Staircases and Related Sequences*

In this section we introduce staircases, which are the main geometric objects we will analyze. Then we will define some related sequences of integers and state basic facts about them and their connection to staircases. We elaborate on the geometric point of view and give additional examples in Section 2.3.

Now, what is a staircase? Let L be an oriented rational line in the plane. Then L defines a positive half-space H. The task is to describe the lattice points in H that are close to L in the sense that they have distance < 1 to the line in the Manhattan metric. Equivalently we consider those points $x \in \mathbb{Z}^2 \cap H$ from which we can reach a point in the other half-space by a single horizontal or vertical step of unit length. Such a set of points we call a staircase. See Figure 2.1 for two examples. Note that it is sufficient to depict the staircase only under a primitive vector generating the line (in the first example the vector $(3, 8)$), as after that (and before that) the same pattern of points is repeated.

We do not require L to pass through the origin (or any other element of $\mathbb{Z}^2$), contrary to what Figure 2.1 might suggest. But we will see later that we can get all the information we want by looking only at lines through the origin. Also, without loss of generality we will restrict our attention to lines with positive slope, as negative slopes will give us, up to mirror symmetry, the same sets.

We will give a precise formulation and proof of this statement in Lemma 2.3.2. Although the proof does not require additional tools and could already be given here, we will pursue the connection between certain sequences and our point sets first, and postpone all observations that are purely concerned with the point sets to Section 2.3. So let's start defining these sets properly, to make it more clear what we are talking about.

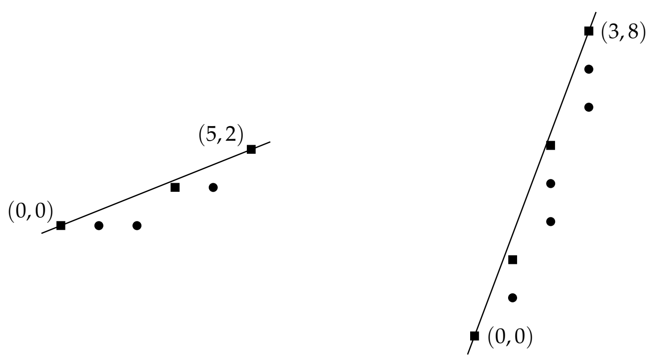

Figure 2.1: A part of the staircases $S_{5,2}$ and $S_{3,8}$. Corners are shown as boxes.

Let $0 < a, b \in \mathbb{N}$ with $\gcd(a,b) = 1$ and let $r \in \mathbb{R}$. These parameters define the line $L_{a,b,r} = \left\{ x \in \mathbb{R}^2 \;\middle|\; x_2 = \frac{b}{a} x_1 + r \right\}$. We denote the closed half-spaces below and above that line by $H^+_{a,b,r}$ and $H^-_{a,b,r}$, respectively. Formally we define for any $\sigma \in \{+1, -1\}$ the **half-space** $H^\sigma_{a,b,r}$ as

$$H^\sigma_{a,b,r} \quad = \quad \left\{ x \in \mathbb{R}^2 \;\middle|\; 0 \le \sigma(\frac{b}{a} x_1 - x_2 + r) \right\}.$$

Most of the time we will use $+$ to represent $+1$ and $-$ to represent -1 and write $\sigma \in \{+, -\}$ for short. Also if $\sigma = +$ and/or $r = 0$ we will omit these parameters and write, e.g., $H_{a,b}$ for $H^{+1}_{a,b,0}$, and similarly for the symbols introduced below. The case $\sigma = +$ and $r = 0$ is of the largest interest to us, as all other cases can be reduced to this one.

The following definitions are illustrated in Figure 2.2. The lattice points in $H^\sigma_{a,b,r}$ that are at distance less than 1 from the line in vertical and horizontal direction, respectively, are

$$\begin{aligned}
\mathcal{V}^\sigma_{a,b,r} &= \mathbb{Z}^2 \cap H^\sigma_{a,b,r} \setminus (H^\sigma_{a,b,r} - \sigma e_2) \\
&= \left\{ z \in \mathbb{Z}^2 \;\middle|\; 0 \le \sigma(\frac{b}{a} z_1 - z_2 + r) < 1 \right\} \\
\mathcal{H}^\sigma_{a,b,r} &= \mathbb{Z}^2 \cap H^\sigma_{a,b,r} \setminus (H^\sigma_{a,b,r} + \sigma e_1) \\
&= \left\{ z \in \mathbb{Z}^2 \;\middle|\; 0 \le \sigma(\frac{b}{a} z_1 - z_2 + r) < \frac{b}{a} \right\}.
\end{aligned}$$

Using this notation we now define the **staircase** $S^\sigma_{a,b,r}$ to be the set of points that are at distance less than 1 from the line in horizontal *or* vertical direction, and we define the **corners** $C^\sigma_{a,b,r}$ to the be the lattice points that are at distance less than 1 in horizontal *and* vertical direction:

$$\begin{aligned}
S^\sigma_{a,b,r} &= \mathcal{V}^\sigma_{a,b,r} \cup \mathcal{H}^\sigma_{a,b,r}, \\
C^\sigma_{a,b,r} &= \mathcal{V}^\sigma_{a,b,r} \cap \mathcal{H}^\sigma_{a,b,r}.
\end{aligned}$$

In other words

$$\begin{aligned}
S^\sigma_{a,b} &= \left\{ z \in \mathbb{Z}^2 \;\middle|\; z \in H^\sigma_{a,b} \text{ but } z - \sigma e_1 \notin H^\sigma_{a,b} \text{ or } z + \sigma e_2 \notin H^\sigma_{a,b} \right\} \\
C^\sigma_{a,b} &= \left\{ z \in \mathbb{Z}^2 \;\middle|\; z \in H^\sigma_{a,b} \text{ but } z - \sigma e_1 \notin H^\sigma_{a,b} \text{ and } z + \sigma e_2 \notin H^\sigma_{a,b} \right\}.
\end{aligned}$$

See Figure 2.1 and also Figure 2.2. Clearly $C_{a,b} \subset S_{a,b}$. For any set $A \subset \mathbb{Z}^2$ and any $x \in \mathbb{Z}$ we call the set $\mathsf{col}_x(A) = \{(x,y) \in A\}$ a **column** of A and for $y \in \mathbb{Z}$ we call the set $\mathsf{row}_y(A) = \{(x,y) \in A\}$ a **row** of A. For any $0 < a, b \in \mathbb{N}$, every row and every column of $S_{a,b}$ contains at least one point and every row and every column of $C_{a,b}$ contains at most one point.

The sequence $(|\mathsf{col}_x(S_{a,b})|)_{x \in \mathbb{Z}}$ is called the **column sequence** of $S_{a,b}$, the sequence $(|\mathsf{row}_y(C_{a,b})|)_{y \in \mathbb{Z}}$ is called the **row sequence** of $C_{a,b}$ and so on.

In the following we summarize some basic facts about staircases. We omit the proofs as they are easy enough to do and would slow the pace of this section without giving the reader further insights. The reader may find it instructive, however, to check the validity of these facts by looking at examples such as those given in the figures of this section.

2.2.1. Fact.

For all $0 < a, b \in \mathbb{N}$ and $\sigma \in \{+, -\}$

$$a \geq b \quad \Leftrightarrow \quad \mathcal{H}^\sigma_{a,b} \subset \mathcal{V}^\sigma_{a,b} \quad \Leftrightarrow \quad \forall x : |\mathsf{col}_x(S^\sigma_{a,b})| = 1 \quad \Leftrightarrow \quad \forall y : |\mathsf{row}_y(C^\sigma_{a,b})| = 1,$$
$$a \leq b \quad \Leftrightarrow \quad \mathcal{H}^\sigma_{a,b} \supset \mathcal{V}^\sigma_{a,b} \quad \Leftrightarrow \quad \forall y : |\mathsf{row}_y(S^\sigma_{a,b})| = 1 \quad \Leftrightarrow \quad \forall x : |\mathsf{col}_x(C^\sigma_{a,b})| = 1.$$

In the former case we call $S^\sigma_{a,b}$ **flat** and in the latter case we call $S^\sigma_{a,b}$ **steep**, see Figure 2.2. Note that this implies $S^\sigma_{a,b} = \mathcal{V}^\sigma_{a,b}$ for flat and $S^\sigma_{a,b} = \mathcal{H}^\sigma_{a,b}$ for steep staircases.

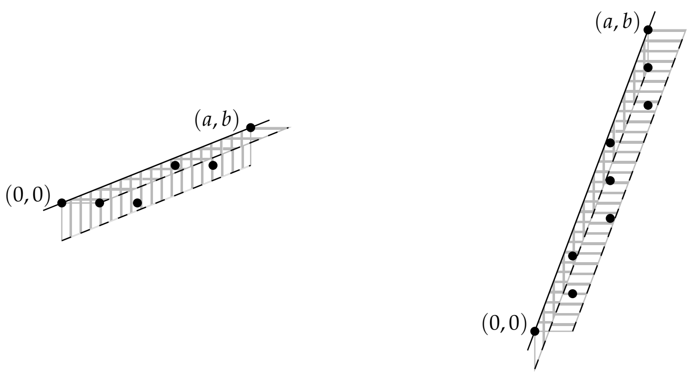

Figure 2.2: This figure shows the flat staircase $S_{5,2}$ and the steep staircase $S_{3,8}$ from Figure 2.1 together with parts of the corresponding sets $\mathcal{H}_{a,b}$ and $\mathcal{V}_{a,b}$. This illustrates Fact 2.2.1.

2.2.2. Fact.

For all $n \in \mathbb{Z}$ the topmost point of $\mathsf{col}_n(S_{a,b})$ is $(n, \left\lfloor \frac{b}{a} n \right\rfloor)$.

If $S_{a,b}$ is flat, we have seen in Fact 2.2.1 that for every $n \in \mathbb{Z}$ the set $\mathsf{col}_n(S_{a,b})$ contains exactly one element, so all elements of $S_{a,b}$ have the form $(n, \left\lfloor \frac{b}{a} n \right\rfloor)$. If $S_{a,b}$ is steep, this is only true for the corners.

This description allows us to compute the difference in height between the topmost points in consecutive columns of $S_{a,b}$. For all $0 < a, b \in \mathbb{N}$ we define the sequence $B_{a,b} = (B_{a,b}(n))_{n \in \mathbb{Z}}$

by

$$B_{a,b}(n) := \left\lfloor \frac{b}{a}n \right\rfloor - \left\lfloor \frac{b}{a}(n-1) \right\rfloor = \frac{b}{a} + \left\{ \frac{b}{a}(n-1) \right\} - \left\{ \frac{b}{a}n \right\}. \tag{2.1}$$

2.2.3. ***Fact.***

If $a \leq b$ (i.e. $S_{a,b}$ is steep) then

$$|\mathsf{col}_n(S_{a,b})| = B_{a,b}(n)$$

and if $b \leq a$ (i.e. $S_{a,b}$ is flat) then

$$|\mathsf{col}_n(C_{a,b})| = B_{a,b}(n),$$

in particular $B_{a,b}$ is a $0,1$-sequence in this case.

The sequence $B_{a,b}$ is the key to connect the geometric description of "points close to a line" with notions from number-theory.

The sequences $(\left\lfloor \frac{b}{a}n \right\rfloor)_{n\in\mathbb{N}}$, $(B_{a,b}(n))_{n\in\mathbb{N}}$ and $(\underline{B}_{a,b}(n))_{n\in\mathbb{N}}$ (see below) are known as the **characteristic sequence**, the **Beatty sequence** and the **Sturmian sequence** of $\frac{b}{a}$, respectively. These sequences are well studied in number theory, see [Sto76, FMT78, PS90, O'B02] for surveys, [Bro93] for historical remarks and [Fra05] for a discussion about the names of the sequences. However, only the characteristic sequences of irrational numbers are non-trivial from the point of view of number theory. They also appear in geometry, see Section 2.8.

For the rest of the section we will establish some definitions connected with the above sequences, and some basic properties of $B_{a,b}$, relating them to the staircases.

In this spirit, instead of working with $(\left\lfloor \frac{b}{a}n \right\rfloor)_{n\in\mathbb{N}}$, $(B_{a,b}(n))_{n\in\mathbb{N}}$ and $(\underline{B}_{a,b}(n))_{n\in\mathbb{N}}$, we will deal with $(\left\lfloor \frac{b}{a}n \right\rfloor)_{n\in\mathbb{Z}}$, $(B_{a,b}(n))_{n\in\mathbb{Z}}$ and $(\underline{B}_{a,b}(n))_{n\in\mathbb{Z}}$, respectively. This is due to the fact that we look at staircases of lines, not of rays.

Now let's define what a Sturmian sequence is. A sequence $s = (s_n)_{n\in\mathbb{Z}}$ of integers $s_n \in \mathbb{Z}$ is called **balanced** (at k) if $s_n \in \{k, k+1\}$ for all $n \in \mathbb{Z}$. If s is balanced at k, we can define a $0,1$-sequence $\underline{s} = (\underline{s}_n)_{n\in\mathbb{Z}}$, which we call the **reduced** sequence, by

$$\underline{s}_n = s_n - k.$$

Note that if $s = (c)_{n\in\mathbb{Z}}$ is constant, s is balanced at both c and $c-1$. In this case $\underline{s}$ is defined with respect to c, i.e. $\underline{s}$ is constant 0.

2.2.4. ***Lemma.***

$B_{a,b}$ is balanced at $\left\lfloor \frac{b}{a} \right\rfloor$. If $\frac{b}{a} \in \mathbb{Z}$, then $B_{a,b}(n) = \frac{b}{a}$ for all $n \in \mathbb{Z}$.

Proof. By (2.1) we know that $|B_{a,b}(n) - \frac{b}{a}| < 1$ and by definition $B_{a,b} \in \mathbb{Z}$. If $\frac{b}{a} \in \mathbb{Z}$, then the fractional parts in (2.1) are both 0, and thus the second statement is also true. □

So Sturmian sequences $(\underline{B}_{a,b}(n))_{n\in\mathbb{N}}$ are well-defined. Furthermore, we now know that only two different integers appear in $B_{a,b}$, and that $\underline{B}_{a,b}$ tells us in which positions the larger integer of the two appears.

Given our geometric interpretation of $B_{a,b}$ from Fact 2.2.3 this means that a steep staircase has columns of only two different lengths and the reduced sequence $(\underline{B}_{a,b}(n))_{n\in\mathbb{N}}$ encodes which columns are long and which columns are short. We will return to the concept of reduction in Section 2.3.

For any sequence $s = (s_n)_{n\in\mathbb{Z}}$ we say that s is **periodic** with period $a \in \mathbb{N}$ if $(s_{n+a})_{n\in\mathbb{Z}} = (s_n)_{n\in\mathbb{Z}}$. We say that a is the **minimal period** of s if there is no period $a' \in \mathbb{N}$ of s with $a' < a$ and write $\mathbb{P}(s)$ for the minimal period of s. By (2.1), if $\gcd(a,b) = 1$, then $B_{a,b}$ is periodic with minimal period a.

For a periodic sequence s we define $\mathsf{period}(s) = (s_n)_{0\leq n<\mathbb{P}(s)}$. If s is a periodic 0, 1-sequence, we write $\mathbb{1}(s)$ for the number of ones in $\mathsf{period}(s)$. We will frequently represent s by the $\mathbb{P}(s)$-tuple $\mathsf{period}(s)$.

As $B_{a,b}$ describes the differences of the maximal heights in adjacent columns of $S_{a,b}$, these differences, accumulated between 0 and $a-1$, must sum up to b.

We summarize the above observations:

2.2.5. Fact.

If $0 < a, b \in \mathbb{N}$ and $\gcd(a,b) = 1$, then

$$\mathbb{P}(B_{a,b}) = a \quad \text{and} \sum_{0\leq n<a} B_{a,b}(n) = b.$$

In particular if $a > b$ (and thus $S_{a,b}$ is flat), then $\mathbb{1}(B_{a,b}) = b$.

To be more flexible when talking about parts of staircases respectively Beatty-sequences, we define the following. Given a sequence $s = (s_n)_{n\in\mathbb{Z}}$, a finite subsequence of the form

$$s|_{[x_0,x_1]} := (s_n)_{x_0\leq n\leq x_1} \text{ for some } x_0 \leq x_1 \in \mathbb{Z}$$

will be called an **interval**. The number of elements $x_1 - x_0 + 1$ of $s|_{[x_0,x_1]}$ we will call the length of the interval and we will denote it by $\mathsf{length}(s|_{[x_0,x_1]})$. If s is a 0, 1-sequence, we will denote the number of ones in an interval $s|_{[x_0,x_1]}$ by $\mathsf{ones}(s|_{[x_0,x_1]})$.

In Fact 2.2.5 we summed over the interval $B_{a,b}|_{[0,a-1]}$. But because of the periodicity of $S_{a,b}$ and $B_{a,b}$ we see that we could have used any interval of length $a-1$. So for any fixed $i \in \mathbb{Z}$ the sequence $B_{a,b}(n+i))_{n\in\mathbb{Z}}$ also describes $S_{a,b}$. This gives rise to the following definition:

We say that sequences $s = (s_n)_{n\in\mathbb{Z}}$ and $s' = (s'_n)_{n\in\mathbb{Z}}$ are identical **up to shift** if there exists an $i \in \mathbb{Z}$ with $(s_{n+i})_{n\in\mathbb{Z}} = (s'_n)_{n\in\mathbb{Z}}$, in symbols $s \equiv s'$. Our goal in Section 2.4 will be to characterize Sturmian sequences up to shift.

2.3. *Geometric Observations*

In this section we develop some properties of staircases and their related sequences from a geometric point of view. The most important operation on staircases is for us the reduction, which we turn to in the latter half of this section. We start with some more elementary operations.

Throughout this section we let $0 < a, b \in \mathbb{N}$ such that $\gcd(a, b) = 1$ and we let $\sigma \in \{+, -\}$.

Elementary Properties of Staircases. As we have already mentioned before, all staircases with a given slope, regardless whether it's the one above or below the line, are translates of each other. Hence they yield the same step sequence up to shift. Before we finally prove this, we state an elementary lemma.

2.3.1. Lemma.

Let $r \in \mathbb{R}$. The line $L_{a,b,r}$ contains a lattice point if and only if $r = \frac{k}{a}$ for some $k \in \mathbb{Z}$.

Proof. Without loss of generality we can assume $-1 < r \leq 0$. For any point $z \in \mathbb{Z}^2$ the vertical distance to the line $L_{a,b}$ is $\frac{b}{a}z_1 - z_2 = \frac{k}{a}$ for some $k \in \mathbb{Z}$. So if $r \neq \frac{k}{a}$ for any $k \in \mathbb{Z}$, then $L_{a,b,r}$ cannot contain a lattice point.

That $r = \frac{k}{a}$ is sufficient for the existence of a lattice point follows directly from the extended Euclidean Algorithm. It can also be shown with this geometric argument: $L_{a,b,r}$ can contain at most one lattice point z with $0 \leq z_1 < a$, for if there were two distinct lattice points with this property then $\gcd(a, b) \neq 1$. On the other hand $\mathcal{V}_{a,b} \cap ([0, a) \times \mathbb{R})$ contains exactly a lattice points, one in each column. Only the lines $L_{a,b,-\frac{k}{a}}$ with $0 \leq k \leq a - 1$ can intersect $\mathcal{V}_{a,b}$. So each of them has to contain at least one lattice point. □

2.3.2. Lemma.

For every $0 < a, b \in \mathbb{N}$ and $r \in \mathbb{R}$

1. $S_{a,b,r} = S_{a,b} + v$ and $C_{a,b,r} = C_{a,b} + v$ for some $v \in \mathbb{Z}^2$ and
2. $S^-_{a,b} = S_{a,b} + v$ and $C^-_{a,b} = C_{a,b} + v$ for some $v \in \mathbb{Z}^2$.

Proof. *1.* By Lemma 2.3.1, if $\frac{k}{a} \leq r < \frac{k+1}{a}$ then $S_{a,b,r} = S_{a,b,\frac{k}{a}}$ and $C_{a,b,r} = C_{a,b,\frac{k}{a}}$. Hence we can assume without loss of generality $r = \frac{k}{a}$, so the line $L_{a,b,r}$ contains a lattice point

$v = (v_1, v_2)$ with $v_2 = \frac{b}{a} v_1 + r$. Then

$$\begin{aligned} z \in H_{a,b,r} - v \quad &\Leftrightarrow \quad z_2 + v_2 \leq \frac{b}{a}(z_1 + v_1) + r \\ &\Leftrightarrow \quad z_2 \leq \frac{b}{a} z_1 \Leftrightarrow z \in H_{a,b}. \end{aligned}$$

This implies the first claim.

2. By Lemma 2.3.1 there is no lattice point v' with $\frac{a-1}{a} < \frac{b}{a} v'_1 - v'_2 < 1$, so

$$\mathcal{V}^-_{a,b} \quad = \quad \left\{ z \in \mathbb{Z}^2 \;\middle|\; 0 \leq -(\frac{b}{a} z_1 - z_2) \leq \frac{a-1}{a} \right\}.$$

Also by Lemma 2.3.1, there exists a lattice point v with $\frac{b}{a} v_1 - v_2 = -\frac{a-1}{a}$ and for this point v

$$\begin{aligned} \mathcal{V}^-_{a,b} - v \quad &= \quad \left\{ z \in \mathbb{Z}^2 \;\middle|\; 0 \leq -(\frac{b}{a}(z_1 + v_1) - (z_2 + v_2)) \leq \frac{a-1}{a} \right\} \\ &= \quad \left\{ z \in \mathbb{Z}^2 \;\middle|\; 0 \leq -(\frac{b}{a} z_1 - z_2 - \frac{a-1}{a}) \leq \frac{a-1}{a} \right\} \\ &= \quad \left\{ z \in \mathbb{Z}^2 \;\middle|\; \frac{a-1}{a} \geq \frac{b}{a} z_1 - z_2 \geq 0 \right\}. \end{aligned}$$

Applying the first observation again, we obtain

$$\mathcal{V}^-_{a,b} - v = \mathcal{V}^+_{a,b}.$$

A similar argument shows $\mathcal{H}^-_{a,b} - v = \mathcal{H}^+_{a,b}$ for a suitable v. Now, because of Fact 2.2.1, $S^\sigma_{a,b} = \mathcal{V}^\sigma_{a,b}$ and $C^\sigma_{a,b} = \mathcal{H}^\sigma_{a,b}$ or $S^\sigma_{a,b} = \mathcal{H}^\sigma_{a,b}$ and $C^\sigma_{a,b} = \mathcal{V}^\sigma_{a,b}$, depending on whether $a > b$ or $a < b$, where $\sigma \in \{+, -\}$. Therefore the above calculations imply 2.3.2.2. □

The previous operations translated the staircases by an integral vector. Now we will introduce some other useful operations. We denote the reflection at the main diagonal by $\nwarrow\!\searrow$ and the reflection at the origin by $\boxtimes$, i.e. we define $\nwarrow\!\searrow(x, y) = (y, x)$ and $\boxtimes(x, y) = (-x, -y)$. Note that both induce involutions, i.e. self-inverse bijective maps, on sets of lattice points; so we understand a set of lattice points if and only if we understand its reflection. The effect of these two reflections on staircases is illustrated with an example in Figures 2.3 and 2.4 and formalized in Lemmas 2.3.3 and 2.3.4.

2.3.3. Lemma.

$$\nwarrow\!\searrow S^\sigma_{a,b} = S^{-\sigma}_{b,a} \text{ and } \nwarrow\!\searrow C^\sigma_{a,b} = C^{-\sigma}_{b,a}.$$

In other words, reflection at the main diagonal swaps numerator and denominator of the slope and places the points on the opposite side of the line. See Figure 2.3.

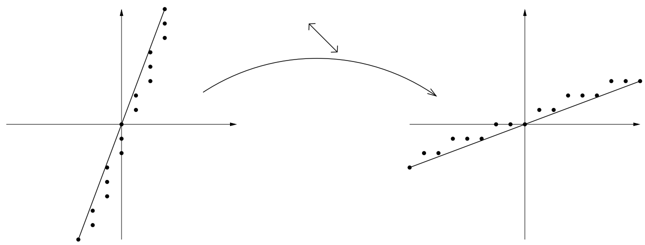

Figure 2.3: The reflection at the main diagonal swaps numerator and denominator of a staircase and places the points on the opposite side of the line. Here we see $\nwarrow\!\!\searrow S_{3,8} = S^{-}_{8,3}$.

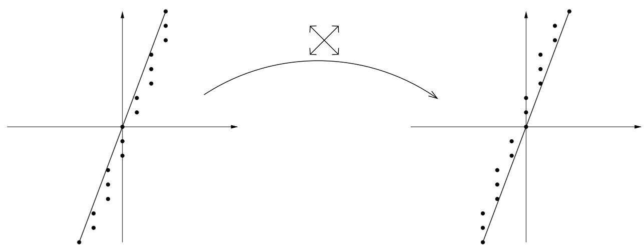

Figure 2.4: The reflection at the origin transforms a staircase below the line into a staircase above the line and vice versa. Here we see $\boxtimes S_{3,8} = S^{-}_{3,8}$. Notice how the column sequence is reversed!

Proof. We compute

$$\begin{aligned}(x,y) \in H^{\sigma}_{a,b} \quad &\Leftrightarrow \quad \sigma y \le \sigma(\frac{b}{a}x) \quad \Leftrightarrow \quad \sigma x \ge \sigma(\frac{a}{b}y) \\ &\Leftrightarrow \quad -(\sigma x) \le -\sigma(\frac{a}{b}y) \quad \Leftrightarrow \quad (y,x) \in H^{-\sigma}_{b,a} \\ &\Leftrightarrow \quad \nwarrow\!\!\searrow(x,y) \in H^{-\sigma}_{b,a}.\end{aligned}$$

This implies both

$$\begin{aligned} z - \sigma e_1 \notin H^{\sigma}_{a,b} \quad &\Leftrightarrow \quad \nwarrow\!\!\searrow(z) + (-\sigma)e_2 \notin H^{-\sigma}_{b,a} \quad \text{and} \\ z + \sigma e_2 \notin H^{\sigma}_{a,b} \quad &\Leftrightarrow \quad \nwarrow\!\!\searrow(z) - (-\sigma)e_1 \notin H^{-\sigma}_{b,a}.\end{aligned}$$

All three equivalences taken together give $\nwarrow\!\!\searrow S^{\sigma}_{a,b} = S^{-\sigma}_{b,a}$ and $\nwarrow\!\!\searrow C^{\sigma}_{a,b} = C^{-\sigma}_{b,a}$. □

2.3.4. *Lemma.*

$\boxtimes S^{\sigma}_{a,b} = S^{-\sigma}_{a,b}$ and $\boxtimes C^{\sigma}_{a,b} = C^{-\sigma}_{a,b}$.

Thus $(|\mathsf{col}_n(S^{\sigma}_{a,b})|)_n = (|\mathsf{col}_{-n}(S^{-\sigma}_{a,b})|)_n$ and $(|\mathsf{col}_n(C^{\sigma}_{a,b})|)_n = (|\mathsf{col}_{-n}(C^{-\sigma}_{a,b})|)_n$.

This means that reflection at the origin maps a staircase below the line to the staircase above the line and vice versa. This operation reverses the Beatty sequence of the staircase. See Figure 2.4.

Proof. We compute

$$\begin{aligned}(x,y) \in H^{\sigma}_{a,b} \quad &\Leftrightarrow \quad \sigma y \leq \sigma(\frac{b}{a}x) \quad \Leftrightarrow \quad -\sigma(-y) \leq -\sigma(\frac{b}{a}(-x)) \\ &\Leftrightarrow \quad (-x,-y) \in H^{-\sigma}_{a,b} \quad \Leftrightarrow \quad \boxtimes(x,y) \in H^{-\sigma}_{a,b}.\end{aligned}$$

which implies both $\boxtimes S^{\sigma}_{a,b} = S^{-\sigma}_{a,b}$ and $\boxtimes C^{\sigma}_{a,b} = C^{-\sigma}_{a,b}$, like in the proof of Lemma 2.3.3.

For the second claim of the lemma we observe (using what we have already shown) that

$$\begin{aligned}(n,y) \in \left(\mathsf{col}_n(S^{\sigma}_{a,b})\right) \quad &\Leftrightarrow \quad (n,y) \in S^{\sigma}_{a,b} \quad \Leftrightarrow \quad (-n,-y) \in \boxtimes S^{\sigma}_{a,b} \\ &\Leftrightarrow \quad (-n,-y) \in S^{-\sigma}_{a,b} \quad \Leftrightarrow \quad (-n,-y) \in \left(\mathsf{col}_{-n}(S^{-\sigma}_{a,b})\right).\end{aligned}$$

As this gives us for any fixed n a bijection between the sets $\mathsf{col}_n(S^{\sigma}_{a,b})$ and $\mathsf{col}_{-n}(S^{-\sigma}_{a,b})$, their cardinality must be the same. The argument for $C^{\sigma}_{a,b}$ is analogous. □

Putting Lemmas 2.3.2.2 and 2.3.4 together, we immediately obtain the non-obvious statement that reversing a Beatty sequence yields the same sequence up to shift.

2.3.5. *Corollary.*

$(|\mathsf{col}_n(S^{\sigma}_{a,b})|)_n \equiv (|\mathsf{col}_{-n}(S^{\sigma}_{a,b})|)_n$.

Proof. $(|\mathsf{col}_n(S^{\sigma}_{a,b})|)_n \equiv (|\mathsf{col}_n(S^{-\sigma}_{a,b}) + v|)_n \equiv (|\mathsf{col}_n(S^{-\sigma}_{a,b})|)_n = (|\mathsf{col}_{-n}(S^{\sigma}_{a,b})|)_n$ □

Similarly, Lemma 2.3.2.1 implies that $C_{a,b}$ and $C_{a,b,r}$ have the same column sequence for any r.

Recursive Description of Staircases. We now return to the operation called reduction, which we defined for balanced sequences in Section 2.2. First, let us observe the relation between Beatty and Sturmian sequences more closely. The following fundamental lemma tells us that, not surprisingly, Sturmian sequences are Beatty sequences with $a > b$ and vice versa.

2.3.6. *Lemma.*

$\underline{B}_{a,b} = B_{a,b \bmod a}$. Conversely if s is a sequence balanced at $k \in \mathbb{N}$ and $\underline{s} = B_{a,b}$, then $s = B_{a,ak+b}$.

Proof. By (2.1) we observe that for any $k \in \mathbb{Z}$ such that both b and $b + ka$ are positive

$$B_{a,b}(n) + k = B_{a,b+ka}(n).$$

$B_{a,b}$ is balanced at $\left\lfloor \frac{b}{a} \right\rfloor = b \;\text{ div }\; a$ by Lemma 2.2.4. Note that by definition $b \bmod a = b - (b \;\text{ div }\; a)\, a$. So

$$\underline{B}_{a,b}(n) = B_{a,b}(n) - b \;\text{ div }\; a = B_{a,b \bmod a}(n).$$

Conversely if $\underline{s}(n) = B_{a,b}(n)$ and s is balanced at k, then

$$s(n) = \underline{s}(n) + k = B_{a,b}(n) + k = B_{a,b+ka}(n).$$

□

How can this relation be phrased in terms of the staircases $S_{a,b}$ and $S_{a,b \bmod a}$? The following lemmas give an answer to this question. See $S_{5,13}$ and $S_{5,3}$ in Figure 2.5.

2.3.7. *Lemma.*

Let $0 < a < b$. The lattice transformation $A = \begin{pmatrix} 1 & 0 \\ b \;\text{ div }\; a & 1 \end{pmatrix}$ gives a bijection between $C_{a,b}$ and $S_{a,b \bmod a}$.

The corners $C_{a,b}$ of the staircase $S_{a,b}$ are just the points of the smaller staircase $S_{a,b \bmod a}$ up to a lattice transform. Here "smaller" refers to both the number of lattice points in a given interval and the encoding length of the two parameters a and b. Note that the inverse of A is $A^{-1} = \begin{pmatrix} 1 & 0 \\ -b \text{ div } a & 1 \end{pmatrix}$.

Proof. As $S_{a,b}$ is steep, $\mathsf{col}_n(C_{a,b}) = \{(n, \left\lfloor \frac{b}{a} n \right\rfloor)\}$ and $\mathsf{col}_n(S_{a,b \bmod a}) = \{(n, \left\lfloor \frac{b \bmod a}{a} n \right\rfloor)\}$ by Facts 2.2.1 and 2.2.2. But

$$\begin{pmatrix} n \\ \left\lfloor \frac{b}{a} n \right\rfloor \end{pmatrix} = \begin{pmatrix} n \\ (b \;\text{ div }\; a) n + \left\lfloor \frac{b \bmod a}{a} n \right\rfloor \end{pmatrix} = A \begin{pmatrix} n \\ \left\lfloor \frac{b \bmod a}{a} n \right\rfloor \end{pmatrix}.$$

□

However, to obtain all points in $S_{a,b}$ from the corners $C_{a,b}$ we need to know which columns of $S_{a,b}$ are long and which are short (in the case $b > a$). It turns out that the corners in long columns are precisely the corners $C_{a,b \bmod a}$ of the smaller staircase, again up to the lattice transformation A.

2.3.8. Lemma.

Let $0 < a < b$. Then $\mathsf{col}_n(C_{a,b \bmod a})$ contains a point if and only if $\mathsf{col}_n(S_{a,b})$ is long.

Proof. $S_{a,b \bmod a}$ is flat. So we know by Fact 2.2.3 that $(n, \left\lfloor \frac{b \bmod a}{a} n \right\rfloor) \in C_{a,b \bmod a}$ if and only if $1 = B_{a,b \bmod a}(n) = \underline{B}_{a,b}(n)$. But this just means that $\mathsf{col}_n(S_{a,b})$ is long. □

Taking the two lemmas together, we can describe every staircase $S_{a,b}$ in terms of the corners and points in the smaller staircase $S_{a,b \bmod a}$. This result, together with our ability to swap the parameters a and b (by means of Lemmas 2.3.3 and 2.3.4) and the fact that the staircases $S_{a,1}$ are easy to describe, we obtain a recursive characterization of all staircases.

Let us look at an example, which is shown in Figure 2.5, before we formulate the recursion formally in Theorem 2.3.9. We want to express $S_{5,13}$ in terms of smaller staircases.

We know that the topmost points in each column are the points in $C_{5,13}$ and $C_{5,13}$ is just the image of $S_{5,3}$ under the lattice transformation $A = \left(\begin{smallmatrix} 1 & 0 \\ 2 & 1 \end{smallmatrix}\right)$. Note that A keeps columns invariant.

We also know that $S_{5,13}$ has columns of lengths 2 and 3 and that the long columns are precisely those in which $C_{5,3}$ contains a point. So if we have an expression for $S_{5,3}$ and $C_{5,3}$, we can give an expression for $S_{5,13}$ and $C_{5,13}$.

To continue this argument inductively, we need to swap the parameters a and b, but this we can achieve by reflecting the staircases at the origin and at the main diagonal. So we reduce the problem of describing $S_{5,3}$ to the problem of describing $S_{3,5}$. We can now continue in this fashion, expressing $S_{3,5}$ in terms of $S_{3,2}$, in terms of $S_{2,3}$, in terms of $S_{2,1}$.

At this point we have finally reached a staircase with integral slope. These staircases have the nice property that all columns and all rows are identical and hence they can be described by a simple expression: the Minkowski sum of the lattice points on a line with those in an interval. This entire process is illustrated in Figure 2.5.

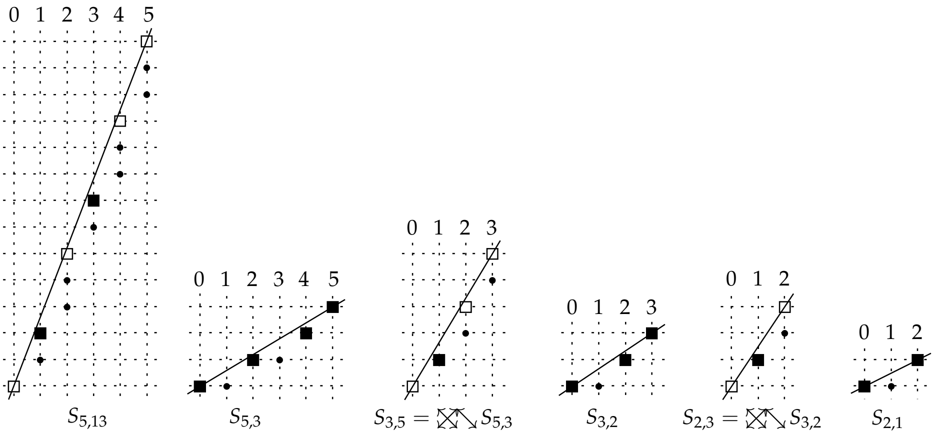

Figure 2.5: This figure show the recursive process of expressing $S_{5,13}$ in terms of smaller staircases, described in the text. In this figure, empty squares indicate the corners of long columns, filled squares corners of short columns. Note that the empty squares occur in $S_{a,b}$ precisely in the columns, in which there is an element of $C_{a,b \bmod a}$.

2.3.9. Theorem.

Let $0 < a, b \in \mathbb{N}$ and $A = \begin{pmatrix} 1 & 0 \\ b \text{ div } a & 1 \end{pmatrix}$.

1. If $a < b$ and $\gcd(a, b) = 1$, then

$$\begin{aligned} C_{a,b} &= AS_{a,b \bmod a} \\ S_{a,b} &= AS_{a,b \bmod a} + \{\left(\begin{smallmatrix}0\\0\end{smallmatrix}\right), \ldots, \left(\begin{smallmatrix}0\\-(b \text{ div } a)+1\end{smallmatrix}\right)\} \\ & \quad \cup AC_{a,b \bmod a} + \{\left(\begin{smallmatrix}0\\-(b \text{ div } a)\end{smallmatrix}\right)\}. \end{aligned}$$

2. $C_{a,b} = \boxtimes\!\nwarrow\!\searrow C_{b,a}$ and $S_{a,b} = \boxtimes\!\nwarrow\!\searrow S_{b,a}$
3. If $b = 1$, then

$$\begin{aligned} C_{a,b} &= \{\left(\begin{smallmatrix}ka\\k\end{smallmatrix}\right) : k \in \mathbb{Z}\} \\ S_{a,b} &= \{\left(\begin{smallmatrix}ka\\k\end{smallmatrix}\right) : k \in \mathbb{Z}\} + \{\left(\begin{smallmatrix}0\\0\end{smallmatrix}\right), \ldots, \left(\begin{smallmatrix}a-1\\0\end{smallmatrix}\right)\} \end{aligned}$$

In Section 2.4 we use this recursive structure to develop a characterization of Sturmian sequences. In Section 2.6 we employ the recursion to obtain short rational functions that enumerate the lattice points inside lattice polytopes in the plane, and in Section 2.7 for a representation of Dedekind-Carlitz polynomials that is computable in polynomial time.

Proof. *1.* By Lemma 2.3.7 the first equation holds. Every column of $S_{a,b}$ contains a corner and every column contains at least $(b \text{ div } a)$ points. So

$$AS_{a,b \bmod a} + \{ \binom{0}{0}, \ldots, \binom{0}{-(b \text{ div } a)+1} \}$$

contains all points in $S_{a,b}$ except the bottom-most points of the long columns. By Lemma 2.3.8 the long columns are precisely those in which $S_{a,b \bmod a}$ has a corner. So $AC_{a,b \bmod a} + \{ \binom{0}{-b \text{ div } a} \}$ is precisely the set of bottom-most points of the long columns of $S_{a,b}$.

2. $S_{a,b} = \boxtimes S^{-}_{a,b} = \boxtimes \nwarrow\!\searrow S_{b,a}$ by Lemmas 2.3.3 and 2.3.4 and similarly for $C_{a,b}$.

3. If $b = 1$, then for all $k, n \in \mathbb{Z}$ we have $\left\lfloor \frac{b}{a} n \right\rfloor = k$ if and only if $ka \leq n \leq (k+1)a - 1$. Hence $\mathsf{row}_k(S_{a,b}) = \{ \binom{ka}{k}, \ldots, \binom{ka+a-1}{k} \}$ and $\mathsf{row}_k(C_{a,b}) = \{ \binom{ka}{k} \}$. □

Relation to the Euclidean Algorithm. This recursion is closely related to the Euclidean Algorithm, which takes as input two natural numbers $c_1, c_2 \in \mathbb{N}$. In each step $c_{i+1} = c_{i-1} \bmod c_i$ is computed. This continues until we reach a j such that $c_{j+1} = 0$ and $c_j \neq 0$. Then $c_j = \gcd(c_1, c_2)$.

Now suppose we want to determine $S_{b,a}$ and $C_{b,a}$ for some $b > a$. We flip the two parameters and then reduce the staircase, i.e. we apply 2.3.9.2 and 2.3.9.1. This reduces the problem to computing $S_{a,b \bmod a}$ and $C_{a,b \bmod a}$. Again we flip and reduce, which reduces the problem to computing $S_{b \bmod a, a \bmod (b \bmod a)}$ and $C_{b \bmod a, a \bmod (b \bmod a)}$ and we continue in this fashion. In other words, we put $c_1 = b$, $c_2 = a$ and $c_{i+1} = c_{i-1} \bmod c_i$ and compute the staircases $S_{c_i, c_{i+1}}$ and $C_{c_i, c_{i+1}}$ recursively, until we arrive at the case $S_{c_{j-1},1}$ and $C_{c_{j-1},1}$ which we can solve directly by 2.3.9.3. That we arrive in this case eventually follows by the correctness of the Euclidean Algorithm and the assumption that $\gcd(a, b) = 1$! Note also that this recursion terminates after few iterations. This is made precise in the following lemma.

2.3.10. Lemma.

Let $a, b \in \mathbb{N}$ and let $(c_n)_{n \in \mathbb{N}}$ denote the sequence defined by $c_1 = b$, $c_2 = a$ and $c_{i+2} = c_i \bmod c_{i+1}$. Then $\min \left\{ j \in \mathbb{N} \;\middle|\; c_{j+1} = 0 \right\} \in \mathcal{O}(\log a)$.

Proof. $(c_i)_{i \geq 2}$ is monotonously decreasing for all $a, b \in \mathbb{N}$, as by definition $c_{i+2} = c_i \bmod c_{i+1} < c_{i+1}$. Thus $c_i \text{ div } c_{i+1} \geq 1$ for $i \geq 2$ and so

$$c_i = (\underbrace{c_i \text{ div } c_{i+1}}_{\geq 1}) c_{i+1} + (\underbrace{c_i \bmod c_{i+1}}_{= c_{i+2}}) \geq \underbrace{c_{i+1}}_{\geq c_{i+2}} + c_{i+2} \geq 2 c_{i+2}$$

for $i \geq 2$. Hence $c_{i+2k} \leq 2^{-k} c_i$, and so if $k \geq \log_2 c_i$, then $c_{i+2k} \leq 1$. In particular the minimal j such that $c_{j+1} = 0$ satisfies $j \leq 2 \log_2 c_2 + 2 \in \mathcal{O}(\log a)$. □

Recursive Description of Parallelepipeds. Instead of describing the infinite set of lattice points in an entire staircase, one might want to describe finite subsets thereof, for example the set of lattice points in only "one period" of the staircase. We now give a recursion for the set of lattice points in the fundamental parallelepipeds of the cones $\mathsf{cone}\left(\binom{a}{b}, \binom{1}{0}\right)$ and $\mathsf{cone}\left(\binom{a}{b}, \binom{0}{-1}\right)$.

The cone generated by $v_1, \ldots, v_n \in \mathbb{R}^d$ is the set

$$\mathsf{cone}(v_1, \ldots, v_n) = \left\{ \sum_{i=1}^{n} \alpha_i v_i \;\middle|\; 0 \leq \alpha_i \in \mathbb{R} \text{ for all } 1 \leq i \leq n \right\}.$$

A cone is rational if all the v_i are rational and it is simplicial if the v_i are linearly independent. The **fundamental parallelepiped** $\Pi_{\mathsf{cone}(v_1,\ldots,v_n)}$ of a simplicial cone $\mathsf{cone}(v_1, \ldots, v_n)$ is defined as

$$\Pi_{\mathsf{cone}(v_1,\ldots,v_n)} := \left\{ \sum_{i=1}^{n} \alpha_i v_i \;\middle|\; 0 \leq \alpha_i < 1 \text{ for all } 1 \leq i \leq n \right\}.$$

Note that any rational cone $\mathsf{cone}(v_1, \ldots, v_n) \subseteq \mathbb{R}^m$ can be transformed unimodularly to a rational cone $\mathsf{cone}(\sigma e_j, v'_1, \ldots, v'_n)$ with $\sigma \in \{+, -\}$ and $1 \leq j \leq m$. So we don't restrict ourselves by considering only cones containing e_1 or $-e_2$ in the generators.

With the above notation

$$\begin{aligned} \mathcal{V}_{a,b} \cap [0, a) \times \mathbb{R} &= \Pi_{\mathsf{cone}\left(\binom{0}{-1}, \binom{a}{b}\right)} \cap \mathbb{Z}^2 \\ \mathcal{H}_{a,b} \cap \mathbb{R} \times [0, b) &= \Pi_{\mathsf{cone}\left(\binom{1}{0}, \binom{a}{b}\right)} \cap \mathbb{Z}^2, \end{aligned}$$

see Figure 2.6. This means that if $a < b$ (and hence $S_{a,b} = \mathcal{H}_{a,b}$), the points z in the staircase $S_{a,b}$ with $0 \leq z_2 < b$ are just the lattice points in the fundamental parallelepiped $\Pi_{\mathsf{cone}\left(\binom{1}{0}, \binom{a}{b}\right)}$. The corners $z \in C_{a,b}$ with $0 \leq z_1 < a$ are just the lattice points in the fundamental parallelepiped $\Pi_{\mathsf{cone}\left(\binom{0}{-1}, \binom{a}{b}\right)}$.

To give an interpretation of our recursion in terms of fundamental parallelepipeds it is convenient to define the set $\Pi^{\circ}_{\mathsf{cone}(v_1,\ldots,v_n)}$ of lattice points (!) in the open fundamental parallelepiped of $\mathsf{cone}(v_1, \ldots, v_n)$ as

$$\Pi^{\circ}_{\mathsf{cone}(v_1,\ldots,v_n)} := \mathbb{Z}^2 \cap \left\{ \sum_{i=1}^{n} \alpha_i v_i \;\middle|\; 0 < \alpha_i < 1 \right\}.$$

Note that if $n = 2$ and both v_1 and v_2 are primitive, then

$$\Pi^{\circ}_{\mathsf{cone}(v_1,v_2)} \cup \left\{ \binom{0}{0} \right\} = \mathbb{Z}^2 \cap \Pi_{\mathsf{cone}(v_1,v_2)}.$$

So it suffices to give a recursion for the sets of lattice points in open fundamental parallelepipeds.

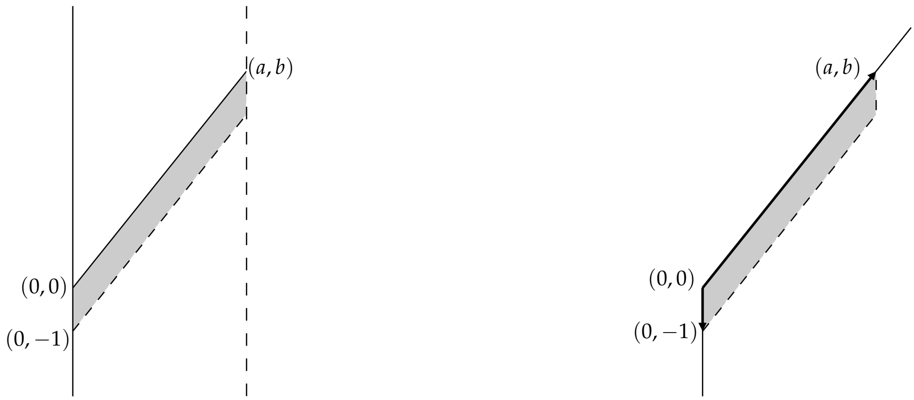

Figure 2.6: If we intersect $\mathcal{V}_{a,b}$ with $[0,a) \times \mathbb{R}$ we obtain the fundamental parallelepiped of the cone generated by $\binom{0}{-1}$ and $\binom{a}{b}$.

We are going to use the following abbreviations:

$$\Pi_{\downarrow,a,b} := \Pi_{\mathsf{cone}(\binom{0}{-1},\binom{a}{b})} \qquad \Pi^\circ_{\downarrow,a,b} := \Pi^\circ_{\mathsf{cone}(\binom{0}{-1},\binom{a}{b})}$$
$$\Pi_{\rightarrow,a,b} := \Pi_{\mathsf{cone}(\binom{1}{0},\binom{a}{b})} \qquad \Pi^\circ_{\rightarrow,a,b} := \Pi^\circ_{\mathsf{cone}(\binom{1}{0},\binom{a}{b})}$$

In terms of open parallelepipeds, Theorem 2.3.9 can now be phrased as follows. An example illustrating the somewhat involved expression in 2.3.11.1 is given in Figure 2.7.

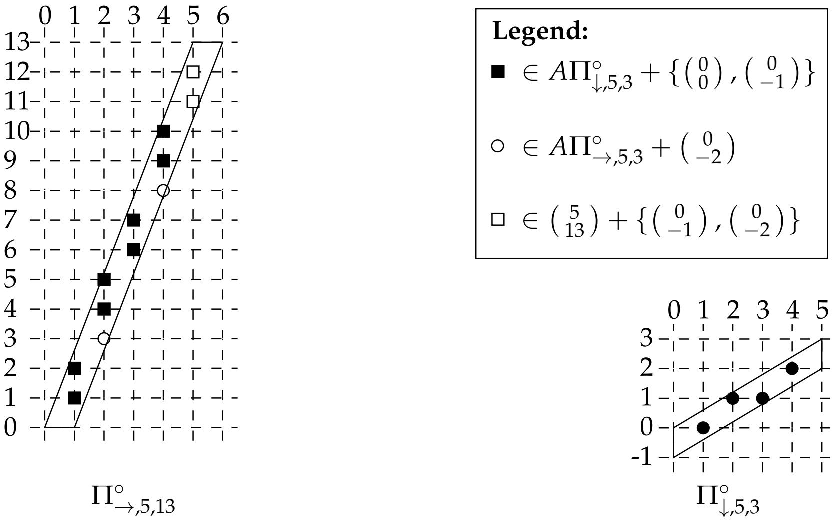

Figure 2.7: This figure illustrates the formula given in 2.3.11.1. $\Pi^\circ_{\rightarrow,5,13}$ is expressed in terms on $\Pi^\circ_{\downarrow,5,3}$ (shown) and $\Pi^\circ_{\rightarrow,5,3}$ (the corners of $\Pi^\circ_{\downarrow,5,3}$). The idea is the same as in Figure 2.5 and Theorem 2.3.9. However there is one important difference: Both $\mathrm{col}_5(\Pi^\circ_{\downarrow,5,3})$ and $\mathrm{col}_5(\Pi^\circ_{\rightarrow,5,3})$ are empty. These have to be added using the third term $\binom{5}{13} + \{\binom{0}{-1}, \binom{0}{-2}\}$.

2.3.11. Theorem.

Let $a, b \in \mathbb{N}$ and $A = \begin{pmatrix} 1 & 0 \\ b \text{ div } a & 1 \end{pmatrix}$.

1. If $0 < a < b$ and $\gcd(a, b) = 1$, then

$$\begin{aligned}
\Pi^\circ_{\downarrow,a,b} &= A\Pi^\circ_{\downarrow,a,b \bmod a} \\
\Pi^\circ_{\rightarrow,a,b} &= A\Pi^\circ_{\downarrow,a,b \bmod a} + \{\tbinom{0}{0}, \ldots, \tbinom{0}{-(b \text{ div } a)+1}\} \\
&\quad \cup A\Pi^\circ_{\rightarrow,a,b \bmod a} + \tbinom{0}{-b \text{ div } a} \\
&\quad \cup \tbinom{a}{b} + \{\tbinom{0}{-1}, \ldots, \tbinom{0}{-b \text{ div } a}\}.
\end{aligned}$$

2. $\Pi^\circ_{\rightarrow,a,b} = \boxtimes\!\nwarrow \Pi^\circ_{\downarrow,b,a} + \binom{a}{b}$ and $\Pi^\circ_{\downarrow,a,b} = \boxtimes\!\nwarrow \Pi^\circ_{\rightarrow,b,a} + \binom{a}{b}$.
3. $\Pi^\circ_{\rightarrow,a,1} = \emptyset$ and $\Pi^\circ_{\downarrow,a,1} = \{\binom{1}{0}, \ldots, \binom{a-1}{0}\}$.

This allows us to describe $\Pi^\circ_{\downarrow,a,b}$ in terms of $\Pi^\circ_{\downarrow,a,b \bmod a}$ and $\Pi^\circ_{\rightarrow,a,b \bmod a}$. The proof is similar to the one of Theorem 2.3.9 and we omit it for brevity.

Recursive Description of Triangles. We conclude this section by giving a similar recursion for triangles, see Figure 2.8. We write

$$\Delta_{a,b} := \mathsf{conv}\left\{ \binom{0}{0}, \binom{a}{0}, \binom{a}{b} \right\}$$

and

$$\Delta'_{a,b} := \Delta_{a,b} \setminus \mathsf{conv}\left\{ \binom{0}{0}, \binom{a}{b} \right\}$$

to denote closed and half-open triangles, respectively. The corresponding lattice point sets are denoted by $T_{a,b} := \Delta_{a,b} \cap \mathbb{Z}^2$ and $T'_{a,b} := \Delta'_{a,b} \cap \mathbb{Z}^2$.

The idea is now that for $0 < a < b$ the triangle $\Delta_{a,b}$ can be decomposed into two parts $\Delta'_{a,(b \text{ div } a)a}$ and $A\Delta_{a,b \bmod a}$. The former is defined by a line with integral slope and hence the set of lattice points $T'_{a,(b \text{ div } a)a}$ is easy to describe. The latter can be transformed into $\Delta_{b \bmod a,a}$ and we can obtain a description of the lattice point set $T_{b \bmod a,a}$ recursively. See Figure 2.8. The resulting recursion is given in Theorem 2.3.12 without proof.

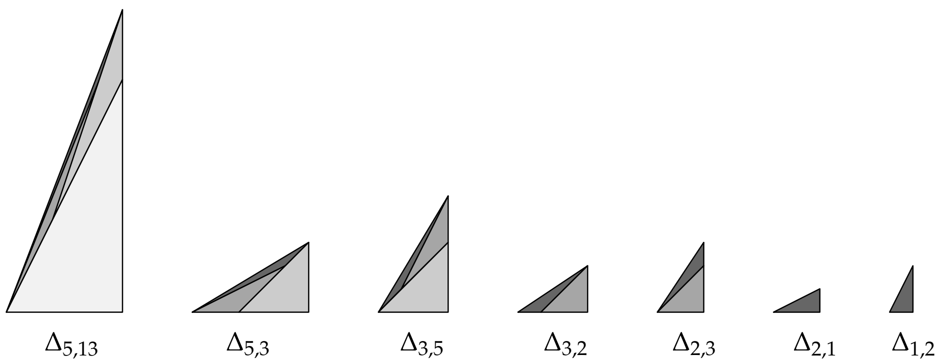

Figure 2.8: Similarly to our recursive description of $S_{5,13}$ (see Figure 2.5), we can apply Theorem 2.3.12 recursively to partition $\Delta_{5,13}$ into triangles with integral slope. The different shadings indicate which triangle the different regions correspond to.

2.3.12. Theorem.

Let $a, b \in \mathbb{N}$ and $A = \begin{pmatrix} 1 & 0 \\ b \text{ div } a & 1 \end{pmatrix}$.

1. $T_{a,b} = AT_{a,b \bmod a} \cup T'_{a,(b \text{ div } a)a}$.
2. $T_{a,b} = \boxtimes\nwarrow\!\!\searrow T_{b,a} + \binom{a}{b}$.
3. $T_{a,1} = \left\{ \binom{k}{0} \mid 0 \leq k \leq a \right\} \cup \left\{ \binom{a}{1} \right\}$.
4. If $k \in \mathbb{N}$, then

$$\begin{aligned} T'_{a,ka} &= \bigcup_{0<l\leq a} \left\{ \binom{l}{0}, \ldots, \binom{l}{lk-1} \right\} \\ &= \left\{ \binom{0}{m} \mid m \in \mathbb{N} \right\} + \left\{ \binom{0}{0}, \binom{1}{0}, \ldots, \binom{a}{0} \right\} \\ &\quad \setminus \left(\left\{ \binom{0}{m} \mid m \in \mathbb{N} \right\} + \left\{ \binom{0}{0}, \binom{1}{k}, \ldots, \binom{a}{ak} \right\} \right). \end{aligned}$$

The advantage of using the second expression for $T'_{a,ka}$ in 2.3.12.4 will become clear in Section 2.6 where we use it to obtain a short rational function representing the generating function of the set of lattice points inside $T'_{a,ka}$. Note that we obtain a recursion formula for T' by replacing every occurrence of T in 2.3.12.1 and 2.3.12.2 with T' and replacing 2.3.12.3 with $T'_{a,1} = \left\{ \binom{k}{0} \mid 1 \leq k \leq a \right\}$.

Kanamaru et al. [KNA94] use a recursive procedure as in Theorem 2.3.12 to give an algorithm to enumerate the set of lattice points *on* a line segment. They go on to give an algorithm that enumerates lattice points inside triangles using the transformation A, however in this case they do not apply recursion and do not mention the partition given in Theorem 2.3.12.1. This partition however is observed by Balza-Gomez et al. [BGMM99]. But as they are interested in giving an algorithm for computing the convex hull of lattice points strictly below a line segment, they do not work with the full set of lattice points $T_{a,b}$. In both cases no explicit recursion formula such as Theorem 2.3.12 is given.

2.4. *Characterizations of Sturmian Sequences*

In this section we state several characterizations of Sturmian sequences of rational numbers, i.e. sequences of the form $\underline{B}_{a,b}$ (or equivalently $B_{a,b}$ with $0 < b \leq a$). We will first motivate each characterization in a separate paragraph without proofs and then summarize them in Theorem 2.4.1. The proof of the theorem occupies Section 2.5.

Recursive Structure. The most important characterization of Sturmian sequences for our purposes is a recursive one. It is based on the concept of reduction presented in Section 2.3 that relates $\underline{B}_{a,b}$ to $\underline{B}_{a,b \bmod a}$ in a way reminiscent of the Euclidean Algorithm.

We first present the idea informally. Let $0 < a < b$ and consider the sequence $B_{b,a}$ and the related staircase $S_{b,a}$. An interval of $B_{b,a}$ of the form $10\ldots0$ of length k corresponds to a corner $c \in S_{b,a}$ and $k-1$ points in $S_{b,a}$ at the same height as c. We call a maximal interval of the form $10\ldots0$ a block. A block of $B_{b,a}$ corresponds to a row of $S_{b,a}$. If the block has length k, the row contains k points. The block sequence $m(B_{b,a})$ of $B_{b,a}$ is the sequence of block lengths of $B_{b,a}$. By the above observation the block sequence of $B_{b,a}$ is the row sequence of $S_{b,a}$ and by Theorem 2.3.9.2 and Corollary 2.3.5 the column sequence of $\boxtimes\!\nwarrow\!\!\searrow S_{b,a} = S_{a,b}$ *up to shift*. This means that $m(B_{b,a}) \equiv B_{a,b}$ is Beatty and hence $\underline{m}(B_{b,a}) \equiv B_{a,b \bmod a}$ is Sturmian. See Figure 2.9 for an example. It turns out that this gives a recursive characterization of Sturmian sequences: a sequence s is Sturmian iff $m(s)$ is balanced and $\underline{m}(s)$ is Sturmian.

To make this precise we give the following definitions. Suppose $s = (s_n)_{n\in\mathbb{Z}}$ is a periodic 0,1-sequence with $\mathbb{P}(s) = a$ and $\mathbb{1}(s) = b$. Without loss of generality we can assume that $s_0 = 1$. Let $i_1, \ldots, i_b \in \{0, \ldots, a-1\}$ be the indices of the 1s, i.e. let $i_1 < i_2 < \ldots < i_b$ with $s_{i_j} = 1$ for all $1 \le j \le b$. Put $i_{b+1} := a$. Then the j-th **block** in $(s_0, \ldots, s_{a-1})$ is $s|_{[i_j, i_{j+1}-1]}$ and the length of the j-th block is $m_j := i_{j+1} - i_j$. The **block sequence** $m(s)$ is the infinite periodic sequence generated by $(m_1, \ldots, m_b)$. Note that this definition determines $m(s)$ only up to shift, which suffices for our purposes. On equivalence classes of sequences up to shift, m is an injective function, i.e. $s_1 \equiv s_2 \Leftrightarrow m(s_1) \equiv m(s_2)$. We call a sequence **block balanced** if it is balanced and its block sequence is balanced. In this case we can consider the reduced block sequence $\underline{m}(s)$ which is again a $0,1$-sequence. A sequence s is **recursively balanced**

- if $\mathbb{1}(s) = 1$, or
- if s is block balanced and $\underline{m}(s)$ is recursively balanced.

The characterization now is this:

A periodic 0,1-sequence is Sturmian if and only if it is recursively balanced.

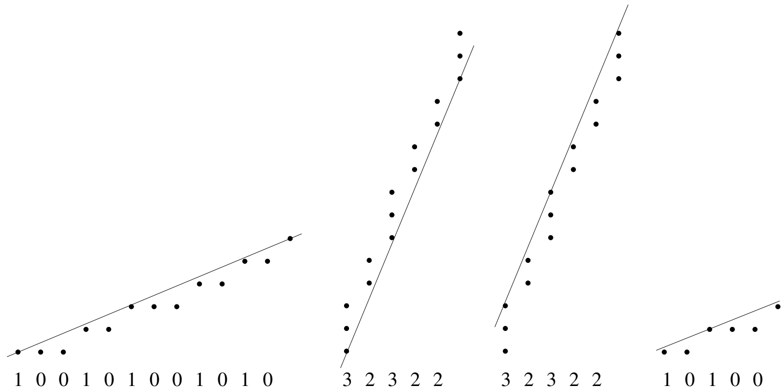

Figure 2.9: *Recursive structure of Sturmian sequences.* This figure shows several staircases $S_{a,b}$ and the corresponding parts of the associated Beatty sequence $B_{a,b}$. In each picture the first and last column correspond to the same element of $B_{a,b}$ modulo the minimal period. The first picture shows $S_{12,5}$ and $B_{12,5}$. $B_{12,5}$ records which columns contain corners, so the block sequence $m(B_{12,5})$ is just the row sequence of $S_{12,5}$. Reflecting at the main diagonal turns rows into columns, so $m(B_{12,5})$ is the column sequence of $\nwarrow\!\!\searrow S_{12,5} = S^{-}_{5,12}$, which is shown in the second picture. Note that in the first picture the part of the staircase that we show was chosen such that points from the first and the last column lie on the defining line - and this property is preserved under reflection at the main diagonal. Applying the reflection at the origin gets us to $\swarrow\!\!\nearrow\!\!\nwarrow\!\!\searrow S_{12,5} = S_{5,12}$, which is shown in the third picture along with its column sequence $B_{5,12}$. The reflection at the origin, however, reverses the column sequence. Fortunately Corollary 2.3.5 tells us that Sturmian sequences are invariant under reversals - up to a shift. This shift in the column sequence can be seen from the fact that in the third picture, the defining line does not pass through points of the first or last column anymore. From the first three pictures, we see that the column sequence of $S_{5,12}$ is just the row sequence of $S_{12,5}$ up to shift, i.e. $m(B_{12,5}) \equiv B_{5,12}$. Now we can apply reduction. We pass from $S_{5,12}$ to $S_{5,2}$, as shown in the last picture, which gives us $\underline{m}(B_{12,5}) \equiv B_{5,2}$. This is a geometric illustration of the combinatorial fact that if s is Sturmian, so is $\underline{m}(s)$.

Even Distribution of 0s and 1s. Common sense suggests that, as the staircase approximates a line, the 0s and 1s of the Sturmian sequence should be distributed as evenly as possible. The actual number of 1s in every interval should be as close as possible to the expected number of 1s. This can be made precise in the following way. On an interval of length l, a line with slope $\frac{b}{a}$ increases by $\frac{b}{a}l$. So the expected number of 1s in an interval of length l of $\underline{B}_{a,b}$ is $\frac{b}{a}l$, if $b < a$, and $\frac{b \bmod a}{a}l$ in general. As $\frac{b \bmod a}{a}l$ is in general not an integer, the best that can be hoped for is that for every interval I of length l the number of 1s contained in l is either $\left\lfloor \frac{b \bmod a}{a}l \right\rfloor$ or $\left\lceil \frac{b \bmod a}{a}l \right\rceil$, and indeed this is a necessary and sufficient characterization of Sturmian sequences. Formally, we say that the 1s in a periodic 0,1-sequence s are **evenly distributed** if for every interval $s|_{[x_0,x_1]}$

$$\mathsf{ones}(s|_{[x_0,x_1]}) \quad \in \quad \left\{ \left\lfloor \frac{\mathbb{1}(s)}{\mathbb{P}(s)} \mathsf{length}(s|_{[x_0,x_1]}) \right\rfloor, \left\lceil \frac{\mathbb{1}(s)}{\mathbb{P}(s)} \mathsf{length}(s|_{[x_0,x_1]}) \right\rceil \right\}. \tag{2.2}$$

Note that if $z \in \mathbb{Z}$ and $r \in \mathbb{R}$, then $z \in \{\lfloor r \rfloor, \lceil r \rceil\}$ if and only if $z - 1 < r < z + 1$, so the condition

$$\mathsf{ones}(s|_{[x_0,x_1]}) - 1 \quad < \quad \tfrac{\mathbb{1}(s)}{\mathbb{P}(s)} \mathsf{length}(s|_{[x_0,x_1]}) \quad < \quad \mathsf{ones}(s|_{[x_0,x_1]}) + 1 \tag{2.3}$$

is equivalent to (2.2). If an interval $s|_{[x_0,x_1]}$ violates the left-hand inequality, then we say it contains too many 1s and if it violates the right-hand inequality, we say it contains too few 1s. The characterization, then, is this.

A periodic $0,1$*-sequence s is Sturmian if and only if the 1s in s are evenly distributed.*

This characterization appears in [GLL78] and was later improved in [Fra05].

Symmetry. A different way to phrase that the 0s and 1s are distributed evenly would be to state that Sturmian sequences are symmetric. If symmetric is taken to mean invariant under reversals (up to shift), then this is a true statement (Corollary 2.3.5) - but insufficient to characterize Sturmian sequences.

However, Lemma 2.3.2 suggests a different notion of symmetry. If we start with a flat staircase $S_{a,b}$ with $0 < b < a$ and move the defining line downwards by a small amount, the resulting staircase $S_{a,b,r}$ will be a translate of $S_{a,b}$. Hence their column sequences are identical up to shift. But using Lemma 2.3.1, we see that if $0 > r \geq -\frac{1}{a}$ the only columns that differ are col_{ka} for $k \in \mathbb{Z}$: the single point in these columns has been moved down by one. The columns col_{ka} do not contain a corner anymore, whereas the columns col_{ka+1} do. In the Sturmian sequence, this translates to taking the corresponding interval $1,0$ and replacing it with the interval $0,1$.

This observation gives rise to the notion of swap symmetry. A periodic $0,1$-sequence s is swap symmetric if there is a pair $(s_i, s_{i+1}) = (1,0)$ such that if we replace this pair and all periodic copies of it by $(0,1)$, we obtain a sequence s' that is identical to s up to a shift. See Figure 2.10 for an example. Formally, given a periodic sequence s and $i \in \mathbb{Z}$, we define the sequence $\mathsf{swap}(s,i) := (\mathsf{swap}(s,i)_n)_n$ by

$$\mathsf{swap}(s,i)_n = \begin{cases} s_n - 1 & \text{if } n \equiv i \mod \mathbb{P}(s) \\ s_n + 1 & \text{if } n \equiv i+1 \mod \mathbb{P}(s) \\ s_n & \text{otherwise} \end{cases}.$$

We call a periodic sequence s **swap symmetric** if there exists an $i \in \mathbb{Z}$ such that $s \equiv \mathsf{swap}(s,i)$. Note that if $s \equiv \mathsf{swap}(s,i)$ then

$$\begin{aligned} s_i &= \mathsf{swap}(s,i)_{i+1} &= s_{i+1} + 1 \\ s_{i+1} &= \mathsf{swap}(s,i)_i &= s_i - 1 \end{aligned}$$

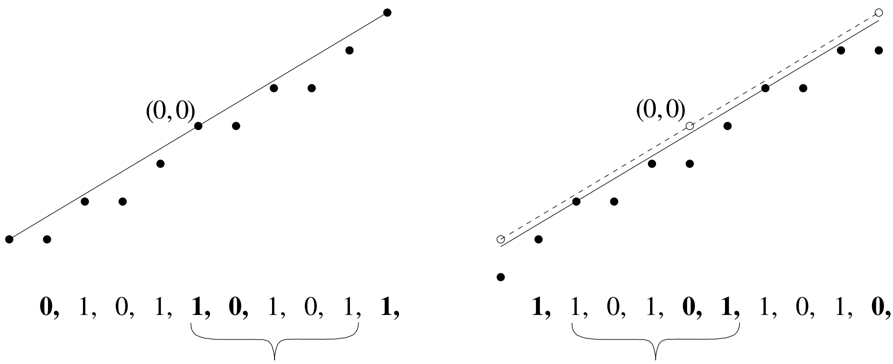

Figure 2.10: *Swap symmetry of Sturmian sequences.* Consider the staircase $S_{5,3}$ and the corresponding Sturmian sequence $B_{5,3}$, which records the columns of $S_{5,3}$ that contain a corner. If we move the line defining $S_{5,3}$ downwards by a small amount, then only those columns change in which there was a point on the line $L_{5,3}$. These columns do not contain corners any more - the corners move one column to the right. So in the sequence $B_{5,3}$ this corresponds to replacing an interval 10 with 01, i.e. by swapping a 1 and a 0. The digits that are swapped are bold in the figure. Now, by Lemma 2.3.2 translating the defining line only shifts the column sequence of $C_{5,3}$. So the sequence obtained from $B_{5,3}$ by swapping is again $B_{5,3}$ up to shift. This is shown by the braces, which indicate minimal periods of both sequences.

as the number of entries $0 \leq k < \mathbb{P}(s)$ such that $s_k = c$ cannot change under a swap for any constant $c \in \mathbb{N}$, for swap symmetric s.

This property characterizes Sturmian sequences:

A periodic 0,1-sequence is Sturmian if and only if it is swap-symmetric.

Having motivated the three characterizations, we can now state the theorem.

2.4.1. Theorem.

Let $s = (s_n)_{n \in \mathbb{Z}}$ be a periodic 0,1-sequence with $\mathbb{P}(s) = a$ and $\mathbb{1}(s) = b \geq 1$. Then the following are equivalent:

(i) $s \equiv \underline{B}_{a,b}$.

(ii) s is recursively balanced.

(iii) The 1s in s are evenly distributed.

(iv) s is swap symmetric.

(i) $\Rightarrow$ *(iii)* is easy to prove (see next section) and *(iii)* $\Rightarrow$ *(i)* was shown by Graham et al. [GLL78], although the concept of "nearly linear" sequences used in [GLL78] differs slightly from 2.4.1*(iii)*. The connection between these definitions[1] was made by Fraenkel [Fra05], where the result from [GLL78] is extended. In both cases the focus lies on the more general case of lines with irrational slope. The proofs given in these two sources differ from the proofs we present in Section 2.5.

As far as we know the concepts of recursively balanced and swap symmetric sequences do not appear in the prior literature.

2.5. *Proof of the Characterizations*

We first show the recursive characterization, i.e. that a periodic 0,1-sequence is Sturmian if and only if it is recursively balanced. To that end we first prove two lemmas.

2.5.1. *Lemma.*

If $0 < a < b$, then $m(B_{b,a}) \equiv B_{a,b}$.

Proof. Let $0 < a < b$. A maximal interval of $B_{b,a}$ of the form $1, 0, \ldots, 0$ corresponds to a row of $S_{b,a}$, so

$$m(B_{b,a}) \equiv (|\mathsf{row}_n(S_{b,a})|)_n = (|\mathsf{col}_n(S^-_{a,b})|)_n \equiv (|\mathsf{col}_n(S_{a,b})|)_n = B_{a,b}$$

where we use Lemma 2.3.3 in the second and Lemma 2.3.2 in the third step. □

From this we also get that for $a > b$ (i.e. flat staircases) the sequence $m(B_{a,b})$ is balanced, and thus Sturmian sequences are block balanced.

2.5.2. *Lemma.*

If s is a block balanced $0,1$-sequence, then s is Sturmian if and only if $\underline{m}(s)$ is Sturmian.

Proof. As s is a block balanced 0,1-sequence

$$\begin{aligned} & s \equiv B_{b,a} \text{ for some } 0 < a < b \\ \Leftrightarrow\quad & m(s) \equiv B_{a,b} \text{ for some } 0 < a < b \\ \Leftrightarrow\quad & \underline{m}(s) \equiv \underline{B}_{a,b} \text{ for some } 0 < a < b \end{aligned}$$

[1] In [Fra05] sequences with an even distribution of 1s are called "balanced".

where the first equivalence holds by Lemma 2.5.1 and the fact that m is injective and the second equivalence holds by Lemma 2.3.6. $\square$

Proof of Theorem 2.4.1: *(i)* $\Leftrightarrow$ *(ii)*. The proof is by induction on $\mathbb{1}(s)$. If $\mathbb{1}(s) = 1$, the statement holds. For the induction step, we have the following equivalences:

$$\begin{aligned} & s \text{ is recursively balanced} \\ \Leftrightarrow\quad & s \text{ is block balanced and} \\ & \underline{m}(s) \text{ is recursively balanced} \\ \Leftrightarrow\quad & s \text{ is block balanced and} \\ & \underline{m}(s) \text{ is Sturmian} \\ \Leftrightarrow\quad & s \text{ is block balanced and} \\ & s \text{ is Sturmian} \\ \Leftrightarrow\quad & s \text{ is Sturmian} \end{aligned}$$

Here we use Lemma 2.5.2 in the third step. Note that the induction terminates in the case $\mathbb{1}(s) = 1$ since if $\mathbb{1}(s) > 1$, then s has blocks of different sizes in a single period and so $\mathbb{1}(s) > \mathbb{1}(\underline{m}(s)) \geq 1$. $\square$

Now we turn to the proof of the characterization that a periodic 0,1-sequence s is Sturmian if and only if the 1s in s are evenly distributed. One direction is easy to show.

Proof of Theorem 2.4.1: *(i)* $\Rightarrow$ *(iii)*. Let $B_{a,b}$ be a Sturmian sequence (i.e. $a > b$) and let $B_{a,b}|_{[x_0,x_1]}$ be any interval of $B_{a,b}$. Using (2.1) and the fact that

$$\frac{b}{a}(x_1 - x_0 + 1) = \frac{\mathbb{1}(B_{a,b})}{\mathbb{P}(B_{a,b})}\mathsf{length}(B_{a,b}|_{[x_0,x_1]})$$

we obtain

$$\mathsf{ones}(B_{a,b}|_{[x_0,x_1]}) = \frac{\mathbb{1}(B_{a,b})}{\mathbb{P}(B_{a,b})}\mathsf{length}(B_{a,b}|_{[x_0,x_1]}) + \left\{\frac{b}{a}(x_0 - 1)\right\} - \left\{\frac{b}{a}x_1\right\}. \tag{2.4}$$

So, since $0 \leq \left\{\frac{b}{a}(x_0 - 1)\right\} < 1$ and $0 \leq \left\{\frac{b}{a}x_1\right\} < 1$ and both terms appear in (2.4) with opposite signs,

$$\mathsf{ones}(B_{a,b}|_{[x_0,x_1]}) - 1 \quad < \frac{\mathbb{1}(B_{a,b})}{\mathbb{P}(B_{a,b})}\mathsf{length}(B_{a,b}|_{[x_0,x_1]}) < \quad \mathsf{ones}(B_{a,b}|_{[x_0,x_1]}) + 1.$$

$\square$

To show the other direction, we make use of the first characterization. We show that if the 1s in s are evenly distributed, then s is recursively balanced.

Proof of Theorem 2.4.1: *(iii)* $\Rightarrow$ *(ii)*. Let s be a periodic 0,1-sequence in which the 1s are evenly distributed. We use induction on $\mathbb{1}(s)$. If $\mathbb{1}(s) = 1$, then by definition s is recursively balanced. For the induction step, we assume $\mathbb{1}(s) > 1$ and show that s is block balanced and the 1s in $\underline{m}(s)$ are evenly distributed. Then we can apply the induction hypothesis to obtain that $\underline{m}(s)$ and hence s is recursively balanced.

Step 1: s is block balanced. If there were blocks of zeros in s that differed in length by at least two, then we could find intervals u and v of the same length l, such that u contains two 1s and v contains none. But then $\{0,2\} \subseteq \{\left\lfloor \frac{\mathbb{1}(s)}{\mathbb{P}(s)} l \right\rfloor, \left\lceil \frac{\mathbb{1}(s)}{\mathbb{P}(s)} l \right\rceil\}$, which is impossible. So s is block balanced at some $k \in \mathbb{N}$, $\underline{m}(s)$ is well defined and the following identities hold.

$$\begin{aligned} \mathbb{P}(\underline{m}(s)) &= \mathbb{1}(s) \\ \mathbb{1}(\underline{m}(s)) &= \mathbb{P}(s) - k\mathbb{1}(s) \end{aligned}$$

Step 2: The 1s in $\underline{m}(s)$ are evenly distributed. Briefly, the idea is this: if m' is an interval of $\underline{m}(s) =: m$ that has too many 1s, looking at the corresponding interval in s we will find many large blocks (i.e. many 0s) and so we can construct an interval s'' of s that has too few 1s. This gives a contradiction to the assumption that the 1s in s are evenly distributed.

Let $m' = \underline{m}(s)|_{[x_0,x_1]}$ be an interval of $\underline{m}(s)$. We have to show that

$$\mathsf{ones}(m') - 1 < \frac{\mathbb{1}(m)}{\mathbb{P}(m)} \mathsf{length}(m') < \mathsf{ones}(m') + 1. \tag{2.5}$$

Assume to the contrary that m' violates (2.5).

We first argue that without loss of generality

$$\mathsf{ones}(m') - 1 \geq \frac{\mathbb{1}(m)}{\mathbb{P}(m)} \mathsf{length}(m'), \tag{2.6}$$

i.e. that m' contains too many 1s. Suppose m' contains too few 1s, i.e. $\mathsf{ones}(m') + 1 \leq \frac{\mathbb{1}(m)}{\mathbb{P}(m)} \mathsf{length}(m')$. Then choose $x_2 > x_1$ such that the length of the interval $m|_{[x_0,x_2]}$ is a multiple $\alpha\mathbb{P}(m)$, $\alpha \in \mathbb{N}$ of the period length. Then $\mathsf{ones}(m|_{[x_0,x_2]}) = \alpha\mathbb{1}(m)$ as m is periodic. Now the interval $m|_{[x_1+1,x_2]}$ has too many 1s, i.e. $\mathsf{ones}(m|_{[x_1+1,x_2]}) - 1 \geq \frac{\mathbb{1}(m)}{\mathbb{P}(m)} \mathsf{length}(m|_{[x_1+1,x_2]})$, as witnessed by the following computation:

$$\begin{aligned} \mathsf{ones}(m|_{[x_1+1,x_2]}) &= \mathsf{ones}(m|_{[x_0,x_2]}) - \mathsf{ones}(m') = \alpha\mathbb{1}(m) - \mathsf{ones}(m') \\ &\geq \alpha\mathbb{1}(m) - \frac{\mathbb{1}(m)}{\mathbb{P}(m)} \mathsf{length}(m') + 1 \\ &= \frac{\mathbb{1}(m)}{\mathbb{P}(m)} (\alpha\mathbb{P}(m) - \mathsf{length}(m')) + 1 \\ &= \frac{\mathbb{1}(m)}{\mathbb{P}(m)} \mathsf{length}(m|_{[x_1+1,x_2]}) + 1. \end{aligned}$$

Each element of m' corresponds to a block of 0s in s, where we take the block to include the preceding 1 but not the succeeding 1. Taking all the blocks in s together that correspond to elements of m' we obtain an interval $s' = s|_{[y_0,y_1]}$ of s. Let $s'' = s|_{[y_0+1,y_1]}$ denote the interval obtained from s' by removing the first 1. Then the following identities hold.

$$\begin{aligned} \mathsf{length}(m') &= \mathsf{ones}(s') = \mathsf{ones}(s'') + 1 \\ \mathsf{ones}(m') &= \mathsf{length}(s') - k\,\mathsf{ones}(s') = \mathsf{length}(s'') + 1 - k(\mathsf{ones}(s'') + 1) \end{aligned}$$

By substituting these and the identities obtained in Step 1 into (2.6) we obtain

$$\mathsf{length}(s'') + 1 - k(\mathsf{ones}(s'') + 1) \quad \geq \quad \frac{\mathbb{P}(s) - k\mathbb{1}(s)}{\mathbb{1}(s)}(\mathsf{ones}(s'') + 1) + 1$$

which by canceling terms implies $\mathsf{length}(s'') \geq \frac{\mathbb{P}(s)}{\mathbb{1}(s)}(\mathsf{ones}(s'') + 1)$ and therefore

$$\mathsf{ones}(s'') + 1 \leq \frac{\mathbb{1}(s)}{\mathbb{P}(s)}\mathsf{length}(s'').$$

This means that s'' is an interval in s with too few 1s, contradicting the assumption that the 1s in s are evenly distributed. □

Finally, we turn to the characterization that a periodic 0,1-sequence is Sturmian if and only if it is swap symmetric. In Section 2.4 we have already tried to motivate that Sturmian sequences are swap symmetric, and the proof indeed proceeds as suggested by Figure 2.10.

Proof of Theorem 2.4.1: *(i) ⇒ (iv)*. Let $0 < b < a$ and let $s = B_{a,b} = (|\mathsf{col}_n(C_{a,b})|)_n$. We claim that

$$\mathsf{swap}(s, 0) = (|\mathsf{col}_n(C_{a,b,-\frac{1}{a}})|)_n \equiv (|\mathsf{col}_n(C_{a,b})|)_n = s$$

which completes the proof. The equivalence in the second step holds by Lemma 2.3.1. All that is left to show is why the first equality holds.

To this end we argue as follows (see Figure 2.10): First we observe that shifting the line down by $-\frac{1}{a}$ only changes those columns $\mathsf{col}_x(S_{a,b})$ with $x \mod a = 0$. More precisely $\mathsf{col}_x(S_{a,b,-\frac{1}{a}}) = \mathsf{col}_x(S_{a,b}) - \binom{0}{1}$ if $0 = x \mod a$ and $\mathsf{col}_x(S_{a,b,-\frac{1}{a}}) = \mathsf{col}_x(S_{a,b})$ otherwise. Now we observe that a point v in a flat staircase is a corner if and only if $v - e_1$ is not in the staircase. As we know which columns changed, and that $a \geq 2$, this allows us to determine where the corners are after the shift. If $x \mod a = 0$, then $|\mathsf{col}_x(C_{a,b,-\frac{1}{a}})| = 0$. If $x \mod a = 1$, then $|\mathsf{col}_x(C_{a,b,-\frac{1}{a}})| = 1$. Otherwise $\mathsf{col}_x(C_{a,b,-\frac{1}{a}}) = \mathsf{col}_x(C_{a,b})$. □

To show that swap symmetric sequences are Sturmian, we first prove two lemmas. Note that in Lemma 2.5.3 we do not claim that if s is swap symmetric, then $m(s)$ is balanced. However we can show that $m(s)$ is swap symmetric.[2] We then observe in Lemma 2.5.4 that swap symmetric sequences are necessarily balanced and hence $m(s)$ is balanced.

[2]Our definition of swap symmetry was phrased such that it can be applied to arbitrary periodic sequences.

2.5.3. Lemma.

Let s be a periodic $0,1$-sequence. If s is swap-symmetric and $\mathbb{1}(s) > 1$, then $m(s)$ is swap-symmetric.

Proof. We know that $\mathbb{P}(m(s)) = \mathbb{1}(s) > 1$. Let $0 \leq i < \mathbb{P}(s)$ be such that $\mathsf{swap}(s,i) \equiv s$. Then $s_i = 1$. Say the block preceding s_i is the j-th block of s. Swapping at i makes all the k-th blocks of s with $k \equiv j \mod \mathbb{1}(s)$ larger by one and all the k-th blocks of s with $k \equiv j+1 \mod \mathbb{1}(s)$ smaller by one while leaving all other blocks unmodified. As $\mathbb{P}(m(s)) > 1$, this means that $\mathsf{swap}(m(\mathsf{swap}(s,i)),j) \equiv m(s)$ but by assumption $\mathsf{swap}(s,i) \equiv s$, so there exists an $l \in \mathbb{Z}$ such that

$$\mathsf{swap}(m(s),l) \equiv \mathsf{swap}(m(\mathsf{swap}(s,i)),j) \equiv m(s)$$

which means that $m(s)$ is swap symmetric. □

2.5.4. Lemma.

If s is a periodic swap-symmetric sequence, then s is balanced.

Proof. Assume to the contrary that s contains at least three different entries. Let $0 \leq i < \mathbb{P}(s)$ be such that $\mathsf{swap}(s,i) \equiv s$. Then $s_i \neq s_{i+1}$. Let $a = s_i$ and $b = s_{i+1}$. Let $c \in \mathbb{Z}$ be such that there exists a $j \in \mathbb{Z}$ with $s_j = c$ but $a \neq c \neq b$. We now define the parameter $d(s)$, which is the sum of the distances of any occurrence of a in $\mathrm{period}(s)$ to the closest preceding occurrence of c in s, i.e.

$$d(s) \quad := \sum_{0 \leq k < \mathbb{P}(s), s_k = a} k - \max\{l < k \mid s_l = c\}.$$

Note that if $s \equiv s'$, then $d(s) = d(s')$ as shifting a sequence to the left or right does not affect the distances between occurrences of values. Now the swap at i interchanges the a at position i and the b at position $i+1$, which increases the distance of this occurrence of a to the previous c by 1 and leaves all other distances of an occurrence of a to a previous c unaffected. Hence $d(\mathsf{swap}(s,i)) = d(s) + 1$ and so $\mathsf{swap}(s,i) \not\equiv s$, which is a contradiction.

So we know that s contains only two different entries and we still have to argue that these entries are consecutive integers. But this is a matter of course as a swap increases one entry by one and decreases another by one. □

After these two lemmas, the proof that swap symmetric sequences are Sturmian is easy. Again we proceed by showing that swap symmetric sequences are recursively balanced.

Proof of Theorem 2.4.1: *(iv) ⇒ (ii).* Let s be a periodic 0,1-sequence that is swap symmetric. There is an index i at which we can swap, so $\mathbb{1}(s) > 0$. If $\mathbb{1}(s) = 1$, s is recursively balanced by definition. So we can assume $\mathbb{1}(s) > 1$. By Lemma 2.5.3 it follows that $m(s)$ is swap symmetric. By Lemma 2.5.4 it follows that $m(s)$ is balanced. Taking both together we conclude that s is block balanced and that $\underline{m}(s)$ is well defined and swap symmetric. By induction we infer that $\underline{m}(s)$ is recursively balanced. But if s is block balanced and $\underline{m}(s)$ is recursively balanced, then s is recursively balanced. $\square$

2.6. *Application: Short Representations*

It is a celebrated result by Barvinok [Bar94] that there is a polynomial time algorithm for counting the number of lattice points inside a given rational polytope when the dimension of the polytope is fixed. Note that if the dimension is an input variable, the problem gets NP-hard [GJ79]. For more about the algorithm see [DL05], [DLHTY04] and the textbook [Bar08].

The crucial ingredient of Barvinok's proof was his result that the set of lattice points in a simplicial cone of any fixed dimension can be expressed using a short generating function. In this section we give a new proof of this result for the special case of 2-dimensional cones (Theorem 2.6.1). A generalization of our proof to higher dimensions is not immediate, we hope, however, that such a generalization can be found in the future.

We consider the Laurent polynomial ring $\mathbb{K}[x_1^{\pm}, \ldots, x_d^{\pm}]$. For a vector $m = (m_1, \ldots, m_d) \in \mathbb{Z}^d$ we write $x^m := x_1^{m_1} \ldots x_d^{m_d}$. This gives a bijection between $\mathbb{Z}^d$ and the set of monomials in $\mathbb{K}[x_1^{\pm}, \ldots, x_d^{\pm}]$. We can thus represent the set of lattice points in a polyhedron P by the generating function $f_P(x) = \sum_{m \in \mathbb{Z}^d \cap P} x^m$. If $P \cap \mathbb{Z}^d$ is large, this representation of f_P contains many terms. Using rational functions it is possible to find shorter representations of f_P. For example the generating function of all non-negative integral multiples of a vector m can be written as $\frac{1}{1-x^m}$, which allows us to express point sets like $\{0, m, \ldots, km\}$ as $\frac{1-x^{(k+1)m}}{1-x^m}$.

Developing these notions in detail is beyond the scope of this thesis. As references we recommend [BR07] and [Bar08]. However, we would like to point out, informally, how the algebraic operations on generating functions correspond to geometric operations: The sum of generating functions corresponds to the union of the respective sets[3]. The product of generating functions corresponds to the Minkowski sum of the respective sets. Taking the product of a generating function and a monomial x^m thus corresponds to translation by m. Evaluating the generating function $f_P(x_1, \ldots, x_d)$ at the values $x^{m_1}, \ldots, x^{m_d}$ for $m_1, \ldots, m_d \in \mathbb{Z}^d$ corresponds to applying the linear map given by the matrix $A = (m_1 \ldots m_d)$ that has the m_i as columns to the set P.

As we already mentioned, for every 2-dimensional rational cone K in $\mathbb{R}^2$ there exists a lattice transform A such that $AK = \operatorname{cone}\left(\binom{1}{0}, \binom{a}{b}\right)$ for $a, b \in \mathbb{N}$ with $\gcd(a, b) = 1$. Barvinok showed a general version of the following theorem for cones of any dimension. We are going to give a new proof of this version for 2-dimensional cones. We make use of the material given at the end of Section 2.3 concerning the recursive description of triangles, in particular the definition of $\Delta'_{a,b}$ and Theorem 2.3.12.

[3] More precisely we are dealing with multisets.

2.6.1. Theorem.

Let $a, b \in \mathbb{N}$ with $\gcd(a,b) = 1$. Let $K = \mathsf{cone}\left(\binom{1}{0}, \binom{a}{b}\right)$. Then f_K admits a representation as a rational function with $\mathcal{O}(\log a)$ terms and this representation can be computed in time polynomial in $\log a + \log b$.

Proof. *Step 1.* We express f_K in terms of $f_{\Delta'_{a,b}}$. To this end we first note that

$$\mathsf{cone}(\binom{1}{0}, \binom{a}{b}) = \bigcup_{k \geq 0} k \binom{a}{b} + (\mathsf{cone}(\binom{1}{0}, \binom{a}{b}) \cap \mathbb{R} \times [0, b))$$

and

$$\begin{aligned} \mathbb{Z}^2 \cap \mathsf{cone}(\binom{1}{0}, \binom{a}{b}) \cap \mathbb{R} \times [0,b) \quad &= \quad \{\binom{0}{0}\} \cup T'_{a,b} \\ &\cup \left(\{\binom{i}{0} \mid i \geq a+1\} + \left\{\binom{0}{j} \,\middle|\, 0 \leq j \leq b-1\right\}\right). \end{aligned}$$

See Figure 2.11. In terms of generating functions this translates into

$$f_K(x) = \frac{1}{1 - x_1^a x_2^b} \left(1 + f_{\Delta'_{a,b}}(x) + \frac{x_1^{a+1}}{1 - x_1} \frac{1 - x_2^b}{1 - x_2}\right).$$

Here we express f_K using $f_{\Delta'_{a,b}}$ and a constant number of other terms. So it suffices to give a short expression of $f_{\Delta'_{a,b}}$.

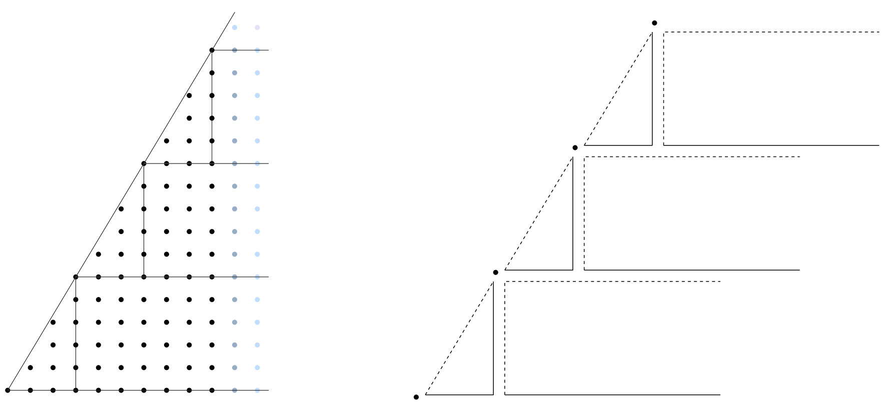

Figure 2.11: Expressing the lattice points in a cone in terms of triangles. The right picture shows the occurring shapes (dashed lines indicate open faces).

Step 2. We use the recursion from Theorem 2.3.12 to give a short expression for $f_{\Delta'_{a,b}}$. Let $(c_n)_n$ be the sequence defined by $c_1 = b$, $c_2 = a$ and $c_{i+2} = c_i \mod c_{i+1}$ and let j be the index such that $c_{j+1} = 1$ and $c_{j+2} = 0$. We express $f_{T'_{c_{i+1},c_i}}$ in terms of $f_{T'_{c_{i+2},c_{i+1}}}$, by applying first 2.3.12.1 and then 2.3.12.2 and 2.3.12.4.

$$\begin{aligned} f_{T'_{c_{i+1},c_i}}(x_1,x_2) &= f_{T'_{c_{i+1},c_{i+2}}}(x_1 x_2^{c_i \text{ div } c_{i+1}}, x_2) \\ &\quad + f_{T'_{c_{i+1},(c_i \text{ div } c_{i+1})c_{i+1}}}(x_1,x_2) \\ &= x_1^{c_{i+1}} x_2^{c_{i+2}} \cdot f_{T'_{c_{i+2},c_{i+1}}}(x_2^{-1}, x_1^{-1} x_2^{-(c_i \text{ div } c_{i+1})}) \\ &\quad + \frac{1}{1-x_2} \cdot \left(\frac{1-x_1^{c_{i+1}+1}}{1-x_1} - \frac{1-x_1^{c_{i+1}+1} x_2^{(c_i \text{ div } c_{i+1})(c_{i+1}+1)}}{1-x_1 x_2^{c_i \text{ div } c_{i+1}}} \right). \end{aligned}$$

We have thus expressed $f_{T'_{c_{i+1},c_i}}$ using a constant number of other terms. We proceed in this fashion until we reach the case $f_{T'_{c_{j+1},c_j}} = f_{T'_{1,c_j}}$ which we can solve directly using 2.3.12.3:

$$f_{T'_{1,c_j}} = x_1 \frac{1-x_2^{c_j}}{1-x_2}.$$

Step 3. The expression is short and can be computed in polynomial time as the Euclidean Algorithm is fast. By Lemma 2.3.10 the number of iterations required in step 2 is $\mathcal{O}(\log a)$. In each step we pick up a constant number of terms. So the total number of terms in the final expression is $\mathcal{O}(\log a)$. The algorithm runs in time polynomial in $\log a + \log b$ as the numbers $c_{i+2} = c_i \mod c_{i+1}$ and $c_i \text{ div } c_{i+1}$ can be computed in time polynomial in $\log c_i + \log c_{i+1}$. □

This proof of Theorem 2.6.1 differs from Barvinok's. Barvinok gives a *signed* decomposition of a cone into *unimodular* cones. We give a *positive* decomposition of the triangle $T'_{a,b}$ into triangles $T'_{c_{i+1},(c_i \text{ div } c_{i+1})c_{i+1}}$ that are not unimodular but easy to describe, i.e. using a constant number of terms.

In this context "positive" means that the 2-dimensional triangle $T'_{a,b}$ is written as a disjoint union of half-open 2-dimensional triangles $T'_{a,b}$. This does not mean that the numerator of the rational function has only positive coefficients. Negative coefficients appear in the "easy" description of the triangles $T'_{c_{i+1},(c_i \text{ div } c_{i+1})c_{i+1}}$.

Theorem 2.3.11 can be used to obtain a short representation of the generating function of the lattice points in the fundamental parallelepiped of any rational cone in the plane. We implement this idea in the proof of Theorem 2.7.1 in Section 2.7. This representation can also be used to give an alternative proof of Theorem 2.6.1. Again the representation is positive in the sense that the set of lattice points in the fundamental parallelepiped is expressed as a disjoint union of Minkowski sums of intervals. But of course still negative coefficients appear as they appear in the representation of intervals. As opposed to the representation based on triangles, the representation based on fundamental parallelepipeds relies on taking products; so with this approach expanding the products in the numerators leads to an expression that is not short any more.

It is also possible to give a recursion similar to 2.3.9, 2.3.11 and 2.3.12 directly for cones. However, in this case the recursion does require us to take differences of sets and we do not obtain a "positive" decomposition. Nonetheless the recursion differs from the one based on the continued fraction expansion of $\frac{b}{a}$ given in [Bar08, Chapter 15].

2.7. *Application: Dedekind-Carlitz Polynomials*

Given $0 < a, b \in \mathbb{N}$ with $\gcd(a, b) = 1$, Carlitz introduced the following polynomial generalization of Dedekind sums, which Beck, Haase and Matthews in [BHM08] call the **Dedekind-Carlitz polynomial**:

$$c_{a,b}(x,y) := \sum_{k=1}^{a-1} x^{k-1} y^{\lfloor \frac{b}{a} k \rfloor}.$$

For a brief overview of the history of and literature about Dedekind sums and the Dedekind-Carlitz polynomial, we refer to [BHM08]. There also the relationship between Dedekind-Carlitz polynomials and the fundamental parallelepipeds of cones (see below) is established. Appealing to Barvinok's Theorem, Beck, Haase and Matthews conclude that the Dedekind-Carlitz polynomial can be computed in polynomial time and must have a short representation,[4] however they do not give such a short representation explicitly. Also they remark that Dedekind sums can be computed efficiently in the style of the Euclidean Algorithm and ask if such a recursive procedure also exists for Dedekind-Carlitz polynomials. In this section we use the recursion for the lattice points inside a fundamental parallelepiped developed in Section 2.3 to give an explicit recursion formula that allows one to compute short representations of Dedekind-Carlitz polynomials in the style of the Euclidean Algorithm.

We first observe that $c_{a,b}$ is the generating function of the set

$$\left\{ z \in \mathbb{Z}^2 \;\middle|\; z_1 = \left\lfloor \frac{b}{a} z_2 \right\rfloor, 1 \leq z_1 \leq a-1 \right\} - \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \Pi^{\circ}_{\downarrow,a,b} - \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

which is just a translate of the set of lattice points in the open fundamental parallelepiped $\Pi^{\circ}_{\downarrow,a,b}$. Hence the recursion given in Theorem 2.3.11 can be used to give a recursion formula for Dedekind-Carlitz sums in the spirit of the Euclidean Algorithm. To this end, we use $d^{\downarrow}_{a,b}(x,y)$ and $d^{\rightarrow}_{a,b}(x,y)$ to denote the generating functions of the sets $\Pi^{\circ}_{\downarrow,a,b}$ and $\Pi^{\circ}_{\rightarrow,a,b}$, respectively. So

$$\begin{aligned} d^{\downarrow}_{a,b}(x,y) &= \sum_{k=1}^{a-1} x^k y^{\lfloor \frac{b}{a} k \rfloor}, \\ d^{\rightarrow}_{a,b}(x,y) &= \sum_{k=1}^{b-1} x^{\lceil \frac{a}{b} k \rceil} y^k. \end{aligned}$$

Now, by simply translating the geometric operations into the language of generating functions, we obtain the following theorem. In [BHM08] this result is derived from Barvinok's Theorem. We give an explicit recursion formula in the proof.

[4] In [BHM08] this is argued even for higher-dimensional Dedekind-Carlitz polynomials.

2.7.1. Theorem.

Let $0 < a, b \in \mathbb{N}$ with $\gcd(a,b) = 1$. Then $c_{a,b}$ admits a representation as a rational function with $\mathcal{O}(\log a)$ terms and this representation can be computed in time polynomial in $\log a + \log b$.

As was said before, the representation we obtain is "positive" in the sense that we build a partition of the set $\Pi^{\circ}_{\downarrow,a,b}$ using Minkowski sums and disjoint unions of intervals. It is not positive in the sense that all coefficients appearing the representation are positive, as the representations of intervals that we use contain coefficients with opposite signs.

It is important to stress that the representation we obtain makes heavy use of Minkowski sums of intervals. In the language of generating functions, this corresponds to taking products of expressions of the form $\frac{1-x^{ku}}{1-x^u}$ for $k \in \mathbb{N}$ and $u \in \mathbb{Z}^2$. Expanding the numerators of these products by applying the distributive law may lead to a numerator with a number of summands exponential in the number of factors of the product. So the expression we obtain is only short, if products are not expanded. We note that this problem does not occur with the representation we used in the proof of Theorem 2.6.1.

Proof. First we note that $c_{a,b}(x,y) = x^{-1} d^{\downarrow}_{a,b}(x,y)$. Now we construct a short representation of $d^{\downarrow}_{a,b}(x,y)$ by applying Theorem 2.3.11 inductively. To that end let $(c_n)_n$ be the sequence defined by $c_1 = b$, $c_2 = a$ and $c_{i+2} = c_i \mod c_{i+1}$ and let j be the index such that $c_{j+1} = 1$ and $c_{j+2} = 0$. Such a j exists because $\gcd(a,b) = 1$.

By Theorem 2.3.11.2 we can assume without loss of generality $c_1 > c_2$. Then for all $i \geq 1$

$$\begin{aligned}
d^{\downarrow}_{c_{i+1},c_i}(x,y) &= d^{\downarrow}_{c_{i+1},c_{i+2}}(xy^{c_i \text{ div } c_{i+1}}, y) \\
d^{\rightarrow}_{c_{i+1},c_i}(x,y) &= \frac{1 - y^{-(c_i \text{ div } c_{i+1})}}{1 - y^{-1}} d^{\downarrow}_{c_{i+1},c_{i+2},(}xy^{c_i \text{ div } c_{i+1}}, y) \\
&\quad + y^{-(c_i \text{ div } c_{i+1})} d^{\rightarrow}_{c_{i+1},c_{i+2},(}xy^{c_i \text{ div } c_{i+1}}, y) \\
&\quad + x^a y^b \frac{y^{-1} - y^{-(c_i \text{ div } c_{i+1})}}{1 - y^{-1}}
\end{aligned}$$

and

$$d^{\rightarrow}_{c_i,c_{i+1}}(x,y) = x^a y^b d^{\downarrow}_{c_{i+1},c_i}(-y,-x) \quad \text{and} \quad d^{\downarrow}_{c_i,c_{i+1}}(x,y) = x^a y^b d^{\rightarrow}_{c_{i+1},c_i}(-y,-x).$$

Together with

$$d^{\rightarrow}_{c_j,c_{j+1}}(x,y) = 0 \quad \text{and} \quad d^{\downarrow}_{c_j,c_{j+1}}(x,y) = \frac{x - x^a}{1 - x}$$

this gives us a recursion formula for $d^{\downarrow}_{a,b}(x,y)$. In each step of the recursion we pick up only a constant number of terms and by Lemma 2.3.10 we need only $\mathcal{O}(\log a)$ steps, so the resulting representation has only $\mathcal{O}(\log a)$ terms. As standard arithmetic operations can be computed in time polynomial in the input length, the algorithm runs in time polynomial in $\log a + \log b$. □

Note that by using products, one can give a representation of the lattice points in an interval of length n with $\mathcal{O}(\log n)$ many terms and *without using rational functions*. Using such a representation, the above proof gives a representation with $\mathcal{O}(\log^2 a)$ terms in time polynomial in $\log a + \log b$. Moreover this representation then is positive in that every coefficient appearing in this expression has a positive sign.

2.8. Application: Theorem of White

To conclude this chapter, we give a partly new proof for a theorem of White [Whi64, pp.390-394], characterizing lattices tetrahedra containing no lattice points but the vertices. Several proofs appeared over the years, e.g. by Noordzij [Noo81], Scarf [Sca85] (based partly on work by R. Howe) and recently Reznick [Rez06], who also gives an overview of the history of this theorem. We construct our proof based on ideas in [Sca85] and mainly [Rez06].

For $(a, b, n) \in \mathbb{Z}^3$ we define the tetrahedron $T_{a,b,n}$ by

$$T_{a,b,n} = \operatorname{conv}\left\{ \begin{pmatrix}0\\0\\0\end{pmatrix}, \begin{pmatrix}1\\0\\0\end{pmatrix}, \begin{pmatrix}0\\1\\0\end{pmatrix}, \begin{pmatrix}a\\b\\n\end{pmatrix} \right\}.$$

This should not be confused with the point set $T_{a,b}$ defined in Section 2.3. As we will not use the latter any more, no ambiguities should arise.

A "hidden" parameter, as Reznick writes, is $c = n - a - b + 1$ (although he considers a slightly different c). We will see that c plays a role equal to the ones of a and b in $T_{a,b,n}$. Also note that $a + b + c = n + 1$.

As we already defined in Section 2.1, two tetrahedra T and T' are equivalent ($T \approx T'$) if there is an affine lattice transformation which takes the vertices of T to the vertices of T'. A lattice simplex T is **clean** if there are no non-vertex lattice points on the boundary. If there are also no lattice points in the interior of T (i.e. the vertices are the only lattice points), then we call T **empty**.

2.8.1. *White's Theorem.*

A lattice tetrahedron T is empty if and only if it is equivalent to $T_{0,0,1}$ or to some $T_{1,d,n}$, where $\gcd(d, n) = 1$ and $1 \leq d \leq n - 1$.

Proof of Theorem 2.8.1, Part "⇐". As we easily see, $T_{0,0,1}$ is empty. So we consider $T_{1,d,n}$, where $\gcd(d, n) = 1$ and $1 \leq d \leq n - 1$. Let $w \in \mathbb{Z}^3 \cap T_{1,d,n}$. As the first coordinate of all vertices of $T_{1,d,n}$ is either 0 or 1, we know $w_1 \in \{0, 1\}$.

If $w_1 = 0$, then $w \in \mathrm{conv}\left\{\begin{pmatrix}0\\0\\0\end{pmatrix}, \begin{pmatrix}0\\1\\0\end{pmatrix}\right\}$ and w is a vertex.

If $w_1 = 1$, then $w \in \mathrm{conv}\left\{\begin{pmatrix}1\\0\\0\end{pmatrix}, \begin{pmatrix}1\\d\\n\end{pmatrix}\right\}$. As $\gcd(d,n) = 1$ the vector $\begin{pmatrix}0\\d\\n\end{pmatrix}$ is primitive and so w is again a vertex.

Therefore $T_{1,d,n}$ and any equivalent tetrahedron is empty. This proves that lattice tetrahedra of the form $T_{0,0,1}$ or $T_{1,d,n}$ with $\gcd(d,n) = 1$ and $1 \leq d \leq n-1$ are necessarily empty. □

To show the sufficiency, we use the following theorem by Reeve. A nice proof can be found in [Rez06, pp.5-6].

2.8.2. Theorem. (Reeve [Ree57])

The lattice tetrahedron T is clean if and only if $T \approx T_{0,0,1}$ or $T \approx T_{a,b,n}$, where

$$n \geq 2, \quad 0 \leq a, b \leq n-1 \quad \text{and}$$
$$\gcd(a,n) = \gcd(b,n) = \gcd(1-a-b,n) = 1.$$

We will now prove the sufficiency to motivate the rest of the section, where we anticipate the results that are stated and proved only afterwards.

Proof of Theorem 2.8.1, Part "⇒". If T is an empty lattice tetrahedron, then in particular it is clean and thus by Theorem 2.8.2 equivalent to $T_{0,0,1}$ or some $T_{a,b,n}$. If $T \approx T_{0,0,1}$, we are done. Otherwise $T \approx T_{a,b,n}$ such that $n \geq 2$, $0 \leq a, b \leq n-1$ and $\gcd(a,n) = \gcd(b,n) = \gcd(1-a-b,n) = 1$. As, by assumption, $T_{a,b,n}$ is empty we can apply first Lemma 2.8.4 and then Lemma 2.8.5 which tells us that one of a, b, c equals 1. Finally, we can see by Lemma 2.8.3 that we can choose the order of the parameters freely, so we get $T \approx T_{1,d,n}$, where $\gcd(d,n) = 1$ and $1 \leq d \leq n-1$. □

It turns out to be useful to describe a clean tetrahedron $T_{a,b,n}$ in a slightly different way:

2.8.3. Lemma.

Let $T_{a,b,n}$ be empty and $0 \leq a, b \leq n-1$. Then

$$T_{a,b,n} \approx \mathrm{conv}\left\{\begin{pmatrix}1\\0\\0\end{pmatrix}, \begin{pmatrix}0\\1\\0\end{pmatrix}, \begin{pmatrix}0\\0\\1\end{pmatrix}, \begin{pmatrix}a\\b\\c\end{pmatrix}\right\}, \quad \text{and } c \geq 1.$$

In the next two proofs we follow mostly [Sca85, pp.411f].

Proof. If $n < a + b$, then

$$\begin{pmatrix}1\\1\\1\end{pmatrix} = \alpha_1 \begin{pmatrix}0\\0\\0\end{pmatrix} + \alpha_2 \begin{pmatrix}1\\0\\0\end{pmatrix} + \alpha_3 \begin{pmatrix}0\\1\\0\end{pmatrix} + \alpha_4 \begin{pmatrix}a\\b\\n\end{pmatrix}, \text{ where}$$

$\alpha_4 = \frac{1}{n}$, $\alpha_3 = 1 - \frac{b}{n}$, $\alpha_2 = 1 - \frac{a}{n}$, and $\alpha_1 = 1 - \alpha_2 - \alpha_3 - \alpha_4 = \frac{a+b-n-1}{n}$. But this means that $0 \leq \alpha_1 < 1$ and $0 < \alpha_2, \alpha_3, \alpha_4 < 1$, and thus $\begin{pmatrix}1\\1\\1\end{pmatrix} \in T_{a,b,n}$, contradicting the assumption that $T_{a,b,n}$ is empty.

So we know that $n \geq a + b$, which proves $c \geq 1$. The affine lattice transformation

$$x \mapsto \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -1 & -1 & 1 \end{pmatrix} x + \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$$

gives us the desired form for $T_{a,b,n}$. □

The key ingredient for the rest of the proof of Theorem 2.8.1 is to look at the Beatty sequences for $\frac{a}{n}$, $\frac{b}{n}$ and $\frac{c}{n}$ simultaneously. For this purpose let us define the sum of the sequences, i.e.

$$\begin{aligned} f(k) &:= B_{n,a}(k) + B_{n,b}(k) + B_{n,c}(k) \\ &= \left\lfloor \frac{a}{n}k \right\rfloor + \left\lfloor \frac{b}{n}k \right\rfloor + \left\lfloor \frac{c}{n}k \right\rfloor - \left\lfloor \frac{a}{n}(k-1) \right\rfloor - \left\lfloor \frac{b}{n}(k-1) \right\rfloor - \left\lfloor \frac{c}{n}(k-1) \right\rfloor. \end{aligned}$$

This function has a strong connection to Theorem 2.8.1 as we see next.

2.8.4. Lemma.

If $T_{a,b,n}$ is empty, then $f(k) = 1$ for $k = 2, \ldots, n-1$.

Proof. Let us first define the function $g(k) := \lceil \frac{a}{n}k \rceil + \lceil \frac{b}{n}k \rceil + \lceil \frac{c}{n}k \rceil$. An easy computation verifies that for a, b, c relatively prime to n and $2 \leq k \leq n-1$, we have

$$f(k) = g(k) - 3 - (g(k-1) - 3) = g(k) - g(k-1).$$

We will now show that $g(k) = k + 2$ for $k = 1, \ldots, n-1$ and also that a, b, c are relatively prime to n.

Suppose that $g(k) \leq k + 1$ for some k. Then we can define a lattice point $\begin{pmatrix}p\\q\\r\end{pmatrix}$ with

$$p \geq \left\lceil \frac{a}{n}k \right\rceil, \quad q \geq \left\lceil \frac{b}{n}k \right\rceil, \quad r \geq \left\lceil \frac{c}{n}k \right\rceil,$$

and $p + q + r = k + 1$. But then we find $\alpha_1, \ldots, \alpha_4$ with

$$\begin{pmatrix} p \\ q \\ r \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & a \\ 0 & 1 & 0 & b \\ 0 & 0 & 1 & c \end{pmatrix} \begin{pmatrix} \alpha_1 \\ \vdots \\ \alpha_4 \end{pmatrix}, \quad \text{where}$$

$\alpha_1 = p - \frac{a}{n}k, \quad \alpha_2 = q - \frac{b}{n}k, \quad \alpha_3 = r - \frac{c}{n}k, \quad \alpha_4 = \frac{1}{n}k$ and

$$\begin{aligned} \sum_{i=1}^{4} \alpha_i &= p + q + r - \frac{a}{n}k - \frac{b}{n}k - \frac{c}{n}k + \frac{k}{n} \\ &= k + 1 - \frac{(a+b+c-1)k}{n} \quad = \quad k+1-k \\ &= 1. \end{aligned}$$

As this means that $\begin{pmatrix} p \\ q \\ r \end{pmatrix}$ is in conv $\left\{ \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} a \\ b \\ c \end{pmatrix} \right\} \approx T_{a,b,n}$, we have a contradiction.

So we know that $g(k) \geq k + 2$.

If one of a, b, c is not relatively prime to n, say it is c, then $\frac{c}{n}k \in \mathbb{Z}$ for some $k \in \{1, \ldots, n-1\}$. So we get

$$\begin{aligned} g(k) + g(n-k) &= \left\lceil \frac{a}{n}k \right\rceil + \left\lceil \frac{b}{n}k \right\rceil + \left\lceil \frac{c}{n}k \right\rceil + \left\lceil a - \frac{a}{n}k \right\rceil + \left\lceil b - \frac{b}{n}k \right\rceil + \left\lceil c - \frac{c}{n}k \right\rceil \\ &= \left\lceil \frac{a}{n}k \right\rceil + \left\lceil \frac{b}{n}k \right\rceil + \left\lceil \frac{c}{n}k \right\rceil + a - \left\lfloor \frac{a}{n}k \right\rfloor + b - \left\lfloor \frac{b}{n}k \right\rfloor + c - \left\lfloor \frac{c}{n}k \right\rfloor, \end{aligned}$$

and by the assumption on c this is at most $a + b + c + 2 = n + 3$. But then either $g(k) \leq k + 1$ or $g(n-k) \leq n - k + 1$, which cannot be true for an empty $T_{a,b,n}$.

If they are all relatively prime to n, then $g(k) + g(n-k) = a + b + c + 3 = n + 4$. Together with $g(k) \geq k + 2$ and $g(n-k) \geq n - k + 2$ we get the desired equality.

We have now seen that $g(k) = k + 2$ for $k = 1, \ldots, n-1$. Together with $f(k) = g(k) - g(k-1)$ this implies $f(k) = 1$ for $k = 2, \ldots, n-1$. □

This is all we need to finish the proof:

2.8.5. Lemma.

If $f(k) = 1$ for $k = 2, \ldots, n-1$, then at least one of a, b or c equals 1.

This is the new part of the proof. It builds onto the observations about Sturmian sequences developed in the first half of this chapter. In particular we make use of the fact that Sturmian sequences are block balanced, and, more generally, that the 1s in a Sturmian sequence are evenly distributed. See Theorem 2.4.1*(iii)*.

Proof. Suppose not. Without loss of generality $c < a, b$. We consider the intervals $B_{n,a}|_{[1,n]}$, $B_{n,b}|_{[1,n]}$ and $B_{n,c}|_{[1,n]}$. As $c \geq 2$, there is another 1 in $B_{n,c}|_{[1,n]}$ apart from $B_{n,c}(n) = 1$. Let m be the position of the 1 preceding $B_{n,c}(n) = 1$, i.e. $1 < m < n$ such that $B_{n,c}(m) = 1$ and $B_{n,c}(k) = 0$ for $m < k < n$. Because $f(k) = 1$ for all $2 \leq k \leq n-1$, we know that $B_{n,a}(m) = B_{n,b}(m) = 0$.

At this point we make a table listing the values of the three intervals at the positions $1, \ldots, n$, filling in the values that we know and marking the values that we have not yet determined with $*$.

$$\begin{array}{rcccccccccc} & 1 & 2 & \ldots & m-1 & m & m+1 & \ldots & n-1 & n \\ B_{n,a}|_{[1,n]} = & 0 & * & \ldots & * & 0 & * & \ldots & * & 1 \\ B_{n,b}|_{[1,n]} = & 0 & * & \ldots & * & 0 & * & \ldots & * & 1 \\ B_{n,c}|_{[1,n]} = & 0 & * & \ldots & * & 1 & 0 & \ldots & 0 & 1 \end{array}$$

Because $f(n-1) = 1$ we know that either $\underline{B}_{n,a}(n-1) = 1$ or $\underline{B}_{n,b}(n-1) = 1$ and we may assume it is $\underline{B}_{n,a}(n-1) = 1$. We are now going to apply the following argument over and over again. By Theorem 2.4.1*(iii)* we know that if we find an interval of length l in a Sturmian sequence that contains t 1s, then any other interval of length l in the same sequence has to contain at least $t-1$ and at most $t+1$ 1s. In this case $\mathsf{ones}(B_{n,a}|_{[n-1,n]}) = 2$ and so both $\mathsf{ones}(B_{n,a}|_{[m-1,m]}) \geq 1$ and $\mathsf{ones}(B_{n,a}|_{[m,m+1]}) \geq 1$, which means $B_{n,a}(m-1) = B_{n,a}(m+1) = 1$ and consequently $B_{n,b}(m-1) = B_{n,b}(m+1) = 0$. Now our table looks as follows.

$$\begin{array}{rcccccccccc} & 1 & 2 & \ldots & m-1 & m & m+1 & \ldots & n-1 & n \\ B_{n,a}|_{[1,n]} = & 0 & * & \ldots * & 1 & 0 & 1 & * \ldots * & 1 & 1 \\ B_{n,b}|_{[1,n]} = & 0 & * & \ldots * & 0 & 0 & 0 & * \ldots * & 0 & 1 \\ B_{n,c}|_{[1,n]} = & 0 & * & \ldots * & 0 & 1 & 0 & \ldots & 0 & 1 \end{array}$$

But now $\mathsf{ones}(B_{n,b}|_{[m-1,m+1]}) = 0$ and so $B_{n,b}(n-2) = 0$ and $B_{n,a}(n-2) = 1$. This gives $\mathsf{ones}(B_{n,a}|_{[n-2,n]}) = 3$ which allows us to deduce $B_{n,a}(m+2) = 1$ and $B_{n,b}(m+2) = 0$. Then we have $\mathsf{ones}(B_{n,b}|_{[m-1,m+2]}) = 0$ and so $B_{n,b}(n-3) = 0$ and $B_{n,a}(n-3) = 1$, which gives $\mathsf{ones}(B_{n,a}|_{[n-3,n]}) = 4$ and so $B_{n,a}(m+3) = 1$ and $B_{n,b}(m+3) = 0$. We continue this argument inductively until we have shown that $B_{n,b}(k) = 0$ for $m+1 \leq k \leq n-1$.

At this point both $B_{n,b}|_{[m-1,n-1]}$ and $B_{n,c}|_{[m+1,n-1]}$ are intervals of consecutive 0s, of length $n-m+1$ and $n-m-1$, respectively, where the latter is maximal. So the blocks of $B_{n,b}$ are strictly larger than the blocks of $B_{n,c}$. As the 1s in Sturmian sequences are evenly distributed, this implies $c = \mathbb{1}(B_{n,c}) > \mathbb{1}(B_{n,b}) = b$, in contradiction to our assumption. □

Chapter 3.

Counting Polynomials and Reciprocity Theorems

The main characters in the next two chapters are five counting polynomials associated with a graph G: the chromatic polynomial χ_G, the integral flow polynomial φ_G, the modular flow polynomial $\bar{\varphi}_G$, the integral tension polynomial θ_G and the modular tension polynomial $\bar{\theta}_G$. The chromatic polynomial was introduced by Birkhoff [Bir12] and Whitney [Whi32a, Whi32b], while the modular flow and tension polynomials were introduced by Tutte in an article [Tut54] on the chromatic polynomial. There the functions φ_G and θ_G are also mentioned implicitly. Much later Kochol [Koc02] realized that φ_G is also a polynomial and his arguments imply that θ_G is a polynomial as well. All of these polynomials are presented in detail in Section 3.2.

The crucial property of each of these is that they are defined as counting functions. For example $\chi_G(k)$ is the number of proper k colorings of the graph G for any positive integer k. But it turns out that each of these functions is actually a polynomial: There is a unique polynomial in one variable that when evaluated at a positive integer k yields $\chi_G(k)$. So we can interpret all five of our counting functions as polynomials and thus we can evaluate them at numbers that are not positive integers. It is natural to ask:

> What does a counting polynomial count *when evaluated at a negative integer*?

Theorems that give an answer to this type of question are known as **combinatorial reciprocity theorems**. We meet several of these in Section 3.3. One of the most famous instances is Stanley's Chromatic Reciprocity Theorem 3.3.1 which gives a combinatorial interpretation of the value of χ_G at negative integers, up to sign. Beck and Zaslavsky [BZ06a] gave a new proof of this theorem, by constructing a geometric object, a so-called inside-out polytope, whose Ehrhart function is the chromatic polynomial of G. They could then use a suitably generalized version of Ehrhart-Macdonald Reciprocity 3.3.2 to derive Stanley's theorem. This approach is particularly elegant as the geometric modeling allows one to *find* the reciprocity theorem instead of just *proving* it: The Chromatic Reciprocity Theorem is easy to prove via induction once the correct claim has been found. Coming up with the correct hypothesis is the hard part and this is where the geometric ansatz is of great help.

In [BZ06b] Beck and Zaslavsky captured the integral flow polynomial as the Ehrhart polynomial of an inside out polytope and applied their method to obtain the Integral Flow Reciprocity Theorem 3.3.3. Independently, Dall [Dal08] and Chen [Che07] achieved the same thing for the integral tension polynomial. The modular tension polynomial is a divisor of the chromatic polynomial, so the Chromatic Reciprocity Theorem 3.3.1 implicitly contains a reciprocity result for the modular tension polynomial, which can be phrased as in the Modular Tension Reciprocity Theorem 3.3.6. So only one case remained open:

> What does the modular flow polynomial count *when evaluated at a negative integer*?

This question was raised by Beck and Zaslavsky in [BZ06b]. We give an answer in the Modular Flow Reciprocity Theorem 3.3.5, which is the main result in this chapter. Independently, Babson and Beck have been working on a similar result; see their unpublished manuscript [BB].

The difficulty in finding a reciprocity theorem for the modular flow polynomial is that there was no inside-out polytope known whose Ehrhart function coincides with $\bar{\varphi}_G$. So the method of Beck and Zaslavsky could not be applied. We came up with a way of modeling $\bar{\varphi}_G$ as the Ehrhart function of a disjoint union of open polytopes. Applying standard Ehrhart-Macdonald Reciprocity then yielded a “geometric” reciprocity statement, which we then had to interpret to obtain a “combinatorial” reciprocity result. The same line of thought can also be applied to obtain a reciprocity theorem for the modular tension polynomial.

Further investigation then showed that in both cases the disjoint unions can be turned into inside-out polytopes by a suitable projection. So we are able to model both the modular flow and tension polynomials as Ehrhart functions in *two* different ways. The benefit of the former construction is a geometric approach to interpreting the values $\bar{\varphi}_G(0)$ and $\bar{\theta}_G(0)$ while the latter construction gives a geometric explanation of the fact that the leading coefficient of both $\bar{\varphi}_G$ and $\bar{\theta}_G$ is 1.

Once a correct interpretation of the value of $\bar{\varphi}_G$ at a negative integer had been found using geometric arguments, we were able to give a purely combinatorial proof of the Modular Flow Reciprocity Theorem 3.3.5.

One of the most important polynomials in graph theory is the Tutte polynomial $T_G(x, y)$. There are a number of ways to define the Tutte polynomial, for instance the recursive definition via a deletion/contraction formula. But even though many values of the Tutte polynomial have a known combinatorial interpretation, there is no definition of the Tutte polynomial as counting function. Thus the question arises:

> What does the Tutte polynomial count?

We cannot expect $T_G(x, y)$ to count something for every value of x, y, simply because T_G takes negative values in all but one of the four quadrants of the plane. But just as we defined our five counting polynomials at all positive integers, we would like to have a unified interpretation of $T_G(x,y)$ as a counting function at all pairs of positive integers $(x,y) \in \mathbb{Z}^2_{\geq 1}$.

It was shown by Tutte that the modular flow and tension polynomials are evaluations of the Tutte polynomial (Theorem 3.11.1). Combining this with the Convolution Formula 3.11.3 by Kook, Reiner and Stanton, leads to an expression of the Tutte polynomial in terms of flows and tensions (cf. Reiner [Rei99, Corollary 2]), but this does not give rise to an interpretation as a *counting* function. Applying the Modular Flow and Tension Reciprocity Theorems, however, *does* lead to a simple expression of the Tutte polynomial in terms of the *reciprocals* of the modular flow and tension polynomials that gives rise to a natural interpretation of the Tutte polynomial as a counting function at all pairs $(x,y) \in \mathbb{Z}^2_{\geq 2}$, see Theorem 3.11.5. This interpretation appears in Reiner [Rei99, Corollary 9], albeit phrased somewhat differently. It is the opinion of the author that this is an important result, which should become more widely known. Combining Theorem 3.11.5 with the interpretations of $\bar{\varphi}_G(0)$ and $\bar{\theta}_G(0)$ allows us to go one step further and interpret the Tutte polynomial as a counting function at all pairs of positive integers $(x,y) \in \mathbb{Z}_{\geq 1}$.

In this way the Convolution Formula together with the Modular Flow and Tension Reciprocity Theorems provide a framework for interpreting the values of $T_G(x,y)$ at *all* pairs of integers in the plane. In the positive quadrant we can interpret $T_G(x,y)$ as a counting function and in the three other quadrants we can interpret $T_G(x,y)$ as a difference of two counting functions. These results are summarized in Theorem 3.11.7.

Again, once the Theorem 3.11.5 is known, we can give a new purely combinatorial proof. Using the equivalence of Theorem 3.11.5 and the Convolution Formula 3.11.3 this also gives a new proof of 3.11.3.

This chapter is organized as follows. In Section 3.1 we give some preliminary definitions, before we introduce colorings, flows and tensions and our five counting polynomials in Section 3.2. In Section 3.3 we present the five reciprocity theorems and our main tool, namely Ehrhart-Macdonald Reciprocity. Before we begin with our proof we need to draw the connection between flows, tensions and linear algebra in Section 3.4 and recall the concept of total unimodularity in Section 3.5. We then have everything we need to give a geometric proof of the Modular Flow Reciprocity Theorem in Section 3.6, which we relate to the theory of inside-out polytopes in Section 3.7. The combinatorial proof is presented in Section 3.8. We then turn to tensions and apply our geometric approach to prove the Modular Tension Reciprocity Theorem in Section 3.9, where we also construct the corresponding inside-out polytope. In Section 3.10 we take up the task of determining the number of components of the disjoint union of polytopes we constructed. This leads interpretations of the values $\bar{\varphi}_G(0)$ and $\bar{\theta}_G(0)$. The Tutte polynomial takes the stage in Section 3.11 where we give the counting interpretation and show its equivalence to the Convolution Formula using the reciprocity theorems. We also present the general framework for interpreting all values of the Tutte polynomial. A purely combinatorial proof of Theorem 3.11.5 concludes the chapter in Section 3.12.

The results in this chapter are joint work with Raman Sanyal.

3.1. *Graphs and Orientations*

The counting functions we will be dealing with this chapter and the next will count objects defined on graphs. We assume that the reader is familiar with graph theory and we will not define every graph theoretic concept that we are going to use. We refer to the textbooks [Wes00] and [Die05] for more information. But as graphs come in many different variations, it is important to say a few words about the notion of a graph that we are going to use.

The graphs we consider are allowed to have loops and/or multiple edges. All graphs are oriented. However we like to think of a graph and its orientation as two separate objects. We implement this as follows: Every graph comes with a fixed but arbitrary orientation σ_1 of its edges. Any other orientation σ_2 can then be encoded by a sign vector that tells us, for each edge e of the graph, whether the orientations of e in σ_1 and σ_2 coincide.

Formally, an (oriented) **graph** G consists of a finite set of vertices V and a finite set of edge E together with two maps $\mathsf{head} : E \to V$ and $\mathsf{tail} : E \to V$. If e is an edge with $\mathsf{head}(e) = v$ and $\mathsf{tail}(e) = u$ for $u, v \in V$, then we say that e is an edge form u to v. e is oriented away from u or is an **out-edge** of u, while e is oriented towards v or is an **in-edge** of v. e is **incident** with both u and v and the two vertices u and v are **adjacent**. We say e is an edge from u to v which we represent graphically by an arrow from u to v. If we just want to denote that u and v are adjacent and do no care about the edge, we write $u \sim v$. A **loop** is an edge e with $\mathsf{head}(e) = \mathsf{tail}(e)$. A **multiple edge** e is an edge such that there exists another edge $e' \neq e$ with $\{\mathsf{head}(e), \mathsf{tail}(e)\} = \{\mathsf{head}(e'), \mathsf{tail}(e')\}$.

Because we want to allow loops and multiple edges, we cannot, as is customary, encode an edge by the pair of vertices it is incident to. However as the head and tail notation is cumbersome, we will use the shorthand $e = uv \in E$ to denote that e is an edge from u to v for vertices $u, v \in V$ that are not required to be distinct. Also we will simply write $G = (V, E)$ to denote that G is a graph with vertex set V and edge set E and suppress the functions head and tail.

An **orientation** of a graph $G = (V, E)$ is a function $\sigma : E \to \{+, -\}$. We denote by ${}_\sigma G$ the graph that arises from G by reversing the orientation of precisely those edges $e \in E$ with $\sigma(e) = -$. Formally, ${}_\sigma G$ is a graph on vertex set V and edge set E where $e \in E$ is an edge from u to v in ${}_\sigma G$ if and only if $\sigma(e) = +$ and e goes from u to v in G or $\sigma(e) = -$ and e goes from v to u in G. We will frequently speak of an orientation as if it were an oriented graph, in which case we are referring to the oriented graph ${}_\sigma G$. Also we will refer to properties of G with respect to an orientation σ, in which case we also mean properties of the graph ${}_\sigma G$.

One could also view the functions σ as reorientations of the oriented graph G. But there is one pitfall that has to be avoided. Consider the graph G with a single vertex and a single edge e that is a loop. Now consider the orientations σ_+ and σ_- given by $e \mapsto +$ and $e \mapsto -$, respectively. The graphs ${}_{\sigma_+}G$ and ${}_{\sigma_-}G$, viewed as 4-tuples $(V, E, \mathsf{head}, \mathsf{tail})$, are identical! So the graph G has only *one* "reorientation", but it has *two* orientations σ_+ and σ_-. The latter behavior is what we want. So let us stress once more: *The graph consisting of a single loop has two orientations.*

This is an important technicality, but now that it has been discussed we can begin to move more quickly and discuss the remaining graph theoretical preliminaries more informally.

By a **walk** in G we mean a walk in the underlying undirected graph, i.e. a walk may traverse an edge in either direction. If we mean a walk that traverses edges e always from $\mathsf{tail}(e)$ to $\mathsf{head}(e)$, we speak of a **directed** or **oriented walk** and similarly for the notions of path and cycle that follow. A walk is **closed** if its first and last vertex coincide.

G is **connected** if there is a walk from any vertex to any other and G is **strongly connected** if there is an oriented walk from any vertex to any other vertex. The **connected components** or simply **components** of G on the one hand and the **strongly connected components** of G on the other hand are defined accordingly. $c(G)$ denotes the number of components of G.

A **path** is a walk that visits every vertex at most once. A single vertex v constitutes a path from v to itself. A **cycle** is a closed walk in which every vertex is visited at most once except the first and last vertex which is visited twice. Note that a loop is a cycle and the only cycle containing a given loop is the loop itself.

For any graph $G = (V, E)$ and any set of edges $S \subset E$ we denote by $G[S]$ the **restriction** of G onto S, by G/S the **contraction** of S and by $G \setminus S$ the **deletion** of S. $G[S]$ is the graph with vertex set V and edge set S where the maps head and tail from G have been restricted to S. $G \setminus S := G[E \setminus S]$. G/S arises from G by contracting all the edges in S in any order. The contraction of an edge identifies the two incident vertices and deletes the edge. The contraction of a loop simply removes the loop. Note that the contraction of a multi edge creates loops.

3.2. *Colorings, Flows and Tensions*

In this section we introduce colorings, flows and tensions along with the corresponding counting polynomials and we gather some well-known facts about them, most of which go back to Tutte [Tut54]. Throughout this section let $G = (V, E)$ denote a graph.

Colorings

Probably the most famous combinatorial objects associated with a graph are its colorings. A k**-coloring** of a graph $G = (V, E)$ is a map $c : V \to \{0, \ldots, k-1\}$. A coloring is called **proper** if $c(u) \neq c(v)$ whenever $u \sim v$. The counting function

$$\chi_G(k) = \#\{c \text{ a proper } k\text{-coloring of } G\}$$

that assigns to each k the number of proper k-colorings of G is called the **chromatic polynomial** of G. That χ_G is in fact a polynomial can be seen, for example, from the fact that χ_G satisfies the following recursion formula.

3.2.1. Theorem. (Tutte [Tut54])

Let $G = (V, E)$ be a graph. If $E = \emptyset$ then $\chi_G(k) = k^{|V|}$. Otherwise let $e \in E$.

1. If e is a bridge, then $\chi_G(k) = (k-1)\chi_{G/e}(k)$.
2. If e is a loop, then $\chi_G(k) = 0$.
3. If e is neither, then $\chi_G(k) = \chi_{G\setminus e}(k) - \chi_{G/e}(k)$.

Such a formula that expresses a counting function f_G determined by a graph G in terms of the corresponding counting functions $f_{G\setminus S}$ and $f_{G/S}$ determined by graphs $G \setminus S$ and G/S, respectively, for some edge set S, we call a **deletion/contraction formula**. In particular, any function f that assigns to any graph G a number $f_G \in \mathbb{Z}$ such that for any $e \in E$

1. if e is a loop of a bridge, then $f_G = f_{G[e]} \cdot f_{G\setminus e}$;
2. if e is neither, then $f_G = af_{G\setminus e} + bf_{G/e}$

for some constants $a, b \in \mathbb{Z}$ is called a **Tutte-Grothendieck invariant**.

Flows

Another classic family of objects associated with a graph are flows. For any abelian group Grp, a Grp-**flow** on G is a map $f : E \to \mathsf{Grp}$ such that at any vertex v we have a conservation of flow: everything that flows into the vertex minus everything that flows out of the vertex equals zero. Formally for every $v \in V$

$$\sum_{e=uv\in E} f(e) - \sum_{e=vu\in E} f(e) = 0. \tag{3.1}$$

Now, there are infinitely many $\mathbb{R}$- or $\mathbb{Z}$-flows on G. So how can we define a counting function in terms of flows? There are two approaches and we consider the classic one first.

If Grp is a finite group, then the number of Grp-flows is clearly finite. We call a Grp-flow **nowhere zero** if $f(e) \neq 0$ for every $e \in E$, in other words if $\mathsf{supp}(f) = E$. Tutte [Tut54] showed that the number of nowhere zero Grp-flows on G depends only on the cardinality of Grp and not on the particular choice of the group Grp. So it makes sense to define $\bar{\varphi}_G(k)$ as the number of nowhere zero Grp-flows for a fixed group Grp with $|\mathsf{Grp}| = k$. In particular we may choose $\mathsf{Grp} = \mathbb{Z}_k$. So, simply put, the function $\bar{\varphi}_G$ defined by

$$\bar{\varphi}_G(k) = \#\{f \text{ a nowhere zero } \mathbb{Z}_k\text{-flow on } G\}$$

is the **modular flow polynomial** or $\mathbb{Z}_k$**-flow polynomial** of G. That this is in fact a polynomial can, again, be seen from a deletion/contraction formula.

3.2.2. Theorem. (Tutte [Tut54])

Let $G = (V, E)$ be a graph. If $E = \emptyset$ then $\bar{\varphi}_G(k) = 1$. Otherwise let $e \in E$.

1. If e is a bridge, then $\bar{\varphi}_G(k) = 0$.
2. If e is a loop, then $\bar{\varphi}_G(k) = (k-1)\bar{\varphi}_{G\setminus e}(k)$.
3. If e is neither, then $\bar{\varphi}_G(k) = \bar{\varphi}_{G/e}(k) - \bar{\varphi}_{G\setminus e}(k)$.

There is another natural way to define a counting function in terms of flows. A k**-flow** f is a $\mathbb{Z}$-flow with $-k < f(e) < k$ for all $e \in E$. We then let $\varphi_G(k)$ count the number of nowhere zero k-flows on G. More precisely, the function φ_G defined by

$$\varphi_G(k) = \#\{f \text{ a nowhere zero } k\text{-flow on } G\}$$

is the **integral flow polynomial** or k**-flow polynomial** of G. Even though the definition of the k-flow polynomial is somewhat more natural than that of the $\mathbb{Z}_k$-flow polynomial, at least in the sense that we are dealing with $\mathbb{Z}$-flows instead of $\mathbb{Z}_k$-flows, it turns out that the k-flow polynomial is not as easy to handle. There is no known deletion/contraction formula for the k-flow polynomial and only recently Kochol [Koc02] was able to show that φ_G is in fact a polynomial.

The good news is that, because φ_G deals with nowhere zero $\mathbb{Z}$-flows, which we can interpret as certain lattice points in the kernel of the incidence matrix of G as we will see in Section 3.4, the k-flow polynomial can be understood geometrically as the Ehrhart function of a so called inside-out polytope. This was done by Beck and Zaslavsky in the articles [BZ06a] and [BZ06b], where the interested reader can also find an interpretation of the chromatic polynomial in terms of inside-out polytopes. We will touch on these ideas in Sections 3.4, 3.6 and 3.7 and we will encounter them again in the next chapter.

Another interesting aspect of $\mathbb{R}$-flows (and $\mathbb{Z}$-flows) is their close relationship to orientations of the underlying graph G. An orientation σ of G is called **totally cyclic** if every edge lies on some directed cycle in ${}_\sigma G$. We define the **sign pattern** $\mathsf{sgn}(f)$ of an $\mathbb{R}$-flow f as the vector $(\mathsf{sgn}(f(e))_{e\in E}$. If f is nowhere zero, then $\mathsf{sgn}(f) \in \{+,-\}^E$.

3.2.3. Theorem. (Greene and Zaslavsky [GZ83])

The totally cyclic orientations of G are precisely the sign patterns of nowhere zero $\mathbb{Z}$-flows on G.

If G is a graph with a bridge e, then any Grp-flow f on G has $f(e) = 0$. In fact, this characterizes when G possesses nowhere zero flows.

3.2.4. Theorem. (Tutte [Tut54])

Let G be a graph. Then the following are equivalent.

1. G has a nowhere zero $\mathbb{Z}$-flow.
2. G has a nowhere zero $\mathbb{R}$-flow.
3. G has a nowhere zero $\mathbb{Z}_k$-flow for some k.
4. G does not have a bridge.

Tensions

We now come to the less widely known concept of tensions. This concept is closely related to the concept of colorings and it is in some sense dual to the concept of flows. A coloring of G with values in $\mathbb{Z}$ is a map $c : V \to \mathbb{Z}$ defined on the vertices of G. Can we associate with c a function $t : E \to \mathbb{Z}$ defined on the edges of G? Perhaps the most natural way of doing this is to define $t(e) = c(v) - c(u)$ for any edge e from u to v. The maps t that arise from colorings in this way are called tensions.

An instructive way to visualize tensions is to picture a drawing of G in the plane where the y-coordinate or height of a vertex v is $c(v)$. The weight $t(e)$ of an edge $e = uv$ in the corresponding tension t is the difference in height between u and v. $t(e)$ is positive if the arrow from u to v points upwards in the drawing and $t(e)$ is negative if the arrow points downwards. If the arrow points sideways, then $t(e) = 0$ and $c(v) = c(u)$. So t is nowhere zero if and only if c is proper. Let us traverse a cycle in G, keeping track of the changes in height: whenever we traverse an edge e in the direction the arrow points we add $t(e)$ to our running total and whenever we traverse an edge e in the opposite direction we subtract $t(e)$ from our running total. As we traverse a cycle we end up at the same height we started. So our running total has to be zero!

This leads to the following characterization of tensions. Let C be any cycle of G, viewed as a set of edges. Let σ be an orientation of $G[C]$ that is totally cyclic, i.e. that turns $G[C]$ into a directed cycle. Then, a map $t : E \to \mathbb{Z}$ is a tension if and only if for any cycle C of G and any totally cyclic orientation of $G[C]$ we have

$$\sum_{\substack{e \in C \\ \sigma(e)=+}} t(e) - \sum_{\substack{e \in C \\ \sigma(e)=-}} t(e) = 0. \tag{3.2}$$

Note that this equation makes sense in any group. So for any abelian group Grp we define a Grp-**tension** to be a map $t : E \to \mathsf{Grp}$ such that (3.2) holds for any cycle C and totally cyclic orientation σ of $G[C]$.

Just as in the flow case we can now construct two counting polynomials. The classic construction is again based on the observation by Tutte that for a finite abelian group Grp the number of nowhere zero Grp-tensions depends only on the cardinality $|\mathsf{Grp}|$ and not on the particular choice of Grp. We thus define the map $\bar{\theta}_G$ given by

$$\bar{\theta}_G(k) = \#\{t \text{ a nowhere zero } \mathbb{Z}_k\text{-tension on } G\}$$

to be the **modular tension polynomial** or $\mathbb{Z}_k$**-tension polynomial** of G. That this is in fact a polynomial again follows from the fact that $\bar{\theta}$ satisfies the following deletion/contraction formula.

3.2.5. Theorem. (Tutte [Tut54])

Let $G = (V, E)$ be a graph. If $E = \emptyset$ then $\bar{\theta}_G(k) = 1$. Otherwise let $e \in E$.

1. If e is a bridge, then $\bar{\theta}_G(k) = (k-1)\bar{\theta}_{G/e}(k)$.
2. If e is a loop, then $\bar{\theta}_G(k) = 0$.
3. If e is neither, then $\bar{\theta}_G(k) = \bar{\theta}_{G\backslash e}(k) - \bar{\theta}_{G/e}(k)$.

On the other hand we can define a k**-tension** to be a $\mathbb{Z}$-tension t with $-k < t(e) < k$ for all edge $e \in E$. The k**-tension polynomial** or simply **tension polynomial** θ_G of G is then given by

$$\theta_G(k) = \#\{t \text{ a nowhere zero } k\text{-tension on } G\}.$$

Again, there is no known deletion/contraction formula for θ and a geometric setup as in Kochol [Koc02] can be used to see that θ_G is in fact a polynomial, see for example Dall [Dal08].

Let us examine the relationship between colorings and tensions more closely. A coloring uniquely determines a tension, but to a given tension there correspond several colorings. Given a coloring c and one integer z_K for every component K of G, the coloring c' defined by $c'(v) = c(v) + z_K$ for every $v \in K$ gives rise to the same tension as c. Conversely, we can fix one vertex v_K for every component of K. Then, a tension t and an integer z_K for every component K of G correspond to a unique coloring with $c(v_K) = z_K$ for all K. This c can be defined as follows. For any $v \in V$ pick a unique path P in G from the v_K in the same component as v to v. Let σ be an orientation of $G[P]$ that turns $G[P]$ into a directed path from v_K to v. Then define

$$c(v) := z_K + \sum_{e \in P, \sigma(e) = +} t(e) - \sum_{e \in P, \sigma(e) = -} t(e),$$

that is we start with z_K and then sum the "oriented weights" of the edges we pass when traversing P from v_K to v. This function c does not depend on the paths P we chose, precisely because of (3.2).

While this does give a bijection between $\mathbb{Z}$-colorings and pairs of $\mathbb{Z}$-tensions and sequences of integers $(z_K)_K$, there is no such bijection between k-colorings and l-tensions. For any k and l there exists a graph that possesses l-tensions that cannot be realized by any k-coloring: Consider a long directed path, whose edges all have weight $t(e) = l - 1$. If the path consists of m edges, the difference between the smallest color and the largest color used in any coloring c corresponding to t would have to be $m(l-1)$. Conversely, if $k > l$ there are colorings that do not induce an l-tension. However, these problems disappear if we pass to the group $\mathbb{Z}_k$.

We can interpret a k-coloring as a map $c : V \to \mathbb{Z}_k$. Then any k-coloring induces a $\mathbb{Z}_k$-tension via $t(e) = c(v) - c(u)$ as described above. And still, the maps $t : E \to \mathbb{Z}_k$ arising in this way are precisely those functions $t : E \to \mathbb{Z}_k$ that satisfy (3.2) for any cycle C and totally cyclic orientation σ of $G[C]$, where (3.2) is now read as an equation in $\mathbb{Z}_k$. Moreover, if we fix one vertex v_K for every component K of G, then given a $\mathbb{Z}_k$-tension t of G and a $z_K \in \mathbb{Z}_k$ for every component K of G, there is a unique k-coloring c with $c(v_K) = z_K$ and $t(e) = c(v) - c(u)$ for every $e = uv \in E$. So for any $\mathbb{Z}_k$-tension there are $k^{c(G)}$ corresponding k-colorings, where $c(G)$ is the number of components of G. Recalling that a $\mathbb{Z}_k$-tension is nowhere zero if and only if the corresponding k-coloring is proper, we obtain the well-known but important fact that the modular tension polynomial is a divisor of the chromatic polynomial.

3.2.6. Theorem. (Tutte [Tut54])

Let G be a graph with c connected components. Then $\chi_G(k) = k^c \cdot \bar{\theta}_G(k)$.

This fact underlines the importance of the not-so-widely-known modular tension polynomial. It can also be shown by comparing the deletion/contraction formulas 3.2.5 and 3.2.1.

We have now explained the relationship between tensions and colorings. But what about the relationship between tensions and flows. We mentioned that they are in some sense dual to each other, but so far we did not give evidence. This duality will become clear when we take a look at these objects from a linear algebra perspective. However, before we turn to linear algebra, we take a look at the relationship between tensions and orientations. This reveals another way in which tensions and flows complement each other.

In some sense the opposite of a totally cyclic orientation is an acyclic orientation. An **acyclic orientation** of a graph G is an orientation σ such that no edge in ${}_\sigma G$ lies on a directed cycle. Again, the sign pattern of a $\mathbb{Z}$-tension t is the vector $\mathsf{sgn}(t) \in \{0, \pm 1\}^E$ and if t is nowhere zero, then $\mathsf{supp}(\mathsf{sgn}(t)) = E$.

3.2.7. Theorem. (Greene and Zaslavsky [GZ83])

The acyclic orientations of G are precisely the sign patterns of nowhere zero $\mathbb{Z}$-tensions on G.

If e is a loop in G, then any Grp-tension on G will necessarily have $t(e) = 0$. Just as in the case of flows, this characterizes which graphs possess nowhere zero tensions.

3.2.8. Theorem. (Tutte [Tut54])

Let G be a graph. Then the following are equivalent.

1. G has a nowhere zero $\mathbb{Z}$-tension.
2. G has a nowhere zero $\mathbb{R}$-tension.
3. G has a nowhere zero $\mathbb{Z}_k$-tension for some k.
4. G does not have a loop.

3.3. *Reciprocity Theorems*

In the last section we have met five counting functions. Let $f : \mathbb{N} \to \mathbb{N}$ be one of them. For any $k \in \mathbb{N}$ the value $f(k)$ was defined as the cardinality of a certain set with a concrete combinatorial interpretation. We then went on to observe that f is actually a polynomial. More precisely, we observed there exists a unique polynomial $p \in \mathbb{K}[k]$ with $f(k) = p(k)$ for all $k \in \mathbb{N}$. We did not and will not make this notational distinction: we denote both f and p by the same symbol f. But now that we have identified f to be a polynomial, we can do something that was not possible before: we can view f as function defined on all of $\mathbb{R}$. In particular we can evaluate f at a negative argument, and ask the question: what does $f(-k)$ mean combinatorially? Combinatorial reciprocity theorems give an answer to this question. Simply put:

> A combinatorial **reciprocity theorem** is a result that gives a combinatorial interpretation of the value of a counting polynomial at a negative integer. In this section we give combinatorial reciprocity theorems for each of the five counting polynomials defined in the previous section.

Perhaps the most famous combinatorial reciprocity theorem is the following result by Stanley that gives a combinatorial interpretation of the value of the chromatic polynomial at a negative integer. We have already met k-colorings and acyclic orientations and we defined the chromatic polynomial $\chi_G(k)$ to count the number of *proper* k-colorings of G. Now we call an arbitrary k-coloring c and an acyclic orientation σ of G **compatible**, if whenever e is oriented from u to v in σ, we have $c(u) \leq c(v)$. If c is proper, then there is a unique acyclic orientation with this property (cf. 3.2.7). Stanley's Reciprocity Theorem, then, states that the chromatic polynomial of G evaluated at $-k$ counts, up to sign, pairs of k-colorings and compatible acyclic orientations of G.

3.3.1. Chromatic Reciprocity Theorem. (Stanley [Sta73])

Let $G = (V, E)$ be a graph. Then for all $0 < k \in \mathbb{N}$

$$(-1)^{|V|} \cdot \chi_G(-k) \quad = \quad \#\{\ (c, \sigma)\ : \begin{array}{l} c \text{ a } k\text{-coloring of } G \\ \sigma \text{ a compatible acyclic orientation of } G\ \}. \end{array}$$

Beck and Zaslavsky gave a beautiful proof of this theorem in [BZ06a]. The essential ingredient of their proof is this remarkable reciprocity theorem about the Ehrhart function of a lattice polytope. Recall that the Ehrhart function L_S of a set $S \subset \mathbb{R}^n$ counts the number of lattice points in kS, i.e. $\mathsf{L}_S(k) = \#\mathbb{Z}^n \cap kS$. If P is a lattice polytope, then L_P is polynomial.

3.3.2. Ehrhart-Macdonald Reciprocity Theorem. (cf. [BR07])

Let P be a lattice polytope. Then

$$(-1)^{\dim P} \mathsf{L}_P(k) = \mathsf{L}_{\operatorname{relint} P}(k).$$

Note that the above is an identity of polynomials and thus it holds in particular for all $k \in \mathbb{Z}$.

Beck and Zaslavsky captured the chromatic polynomial of a graph G as the Ehrhart function of a so called inside-out polytope. A variant of Ehrhart-Macdonald Reciprocity, suitably generalized to inside-out polytopes, then yielded a natural of proof Stanley's Reciprocity Theorem. This close relation between these two reciprocity theorems is no coincidence. Indeed, Ehrhart-Macdonald Reciprocity is not only a remarkably elegant result, one can also derive reciprocity theorems for all five counting polynomials introduced in the last section from Ehrhart-Macdonald Reciprocity.

This is particularly appealing as this method allows the user to *find* reciprocity theorems: Once the correct statement of a reciprocity theorem has been found, it is usually not too difficult to give an inductive proof of the theorem, often using a deletion/contraction formula. Taking the geometric approach, however, it is often possible to come up with the correct claim while doing the proof. We will see one instance of this in Section 3.6.

So, put somewhat polemically, the central theme of this chapter is this.

> *Ehrhart-Macdonald Reciprocity is the Right Way™ to prove combinatorial reciprocity theorems.*

After this advertisement, we return to the business at hand and formulate the reciprocity theorems for the remaining four counting polynomials, as promised at the beginning of this section. We begin with the k-flow and the k-tension polynomial.

Two vectors $v, w \in \mathbb{R}^n$ are **sign-compatible** if there does not exist an index i such that $\mathsf{sgn}(v_i) = -\mathsf{sgn}(w_i) \neq 0$. For the purpose of the Integral Flow Reciprocity Theorem, we call a k-flow f on G and a totally cyclic orientation σ of G **compatible** if they are sign compatible. One way of interpreting this would be to say that the flow of f along any edge e must be in the direction prescribed by σ. By 3.2.3, we have that if f is nowhere zero, then there is a unique totally cyclic orientation of G that is compatible with f.

3.3.3. Integral Flow Reciprocity Theorem. (Beck and Zaslavsky [BZ06b])

Let $G = (V, E)$ be a graph with c components. Then for all $0 < k \in \mathbb{N}$

$$(-1)^{|E|-|V|+c} \cdot \varphi_G(-k) = \#\{\ (f, \sigma)\ :\ f \text{ a } k\text{-flow on } G,\ \sigma \text{ a compatible totally cyclic orientation of } G\ \}.$$

This reciprocity result for k-flows was shown by Beck and Zaslavsky [BZ06b], using their framework of inside-out polytopes.

The corresponding result for k-tensions was shown by Chen [Che07] and Dall [Dal08], independently. In the context of the Integral Tension Reciprocity Theorem, a k-tension t on G and an acyclic orientation σ of G are **compatible** if they are sign compatible. By 3.2.7, if t is a nowhere zero tension then there is a unique acyclic orientation that is compatible with t.

3.3.4. Integral Tension Reciprocity Theorem. (Chen [Che07] and Dall [Dal08])

Let $G = (V, E)$ be a graph with c components. Then for all $0 < k \in \mathbb{N}$

$$(-1)^{|V|-c} \cdot \theta_G(-k) \quad = \quad \#\{\ (t,\sigma)\ : \begin{array}{l} t \text{ a } k\text{-tension on } G \\ \sigma \text{ a compatible acyclic orientation of } G\ \}. \end{array}$$

So we know reciprocity theorems for the integral flow and tension polynomials. How about the modular flow and tension polynomials?

In a way, we already have a reciprocity theorem for the modular tension polynomial. Theorem 3.2.6 tells us that the $\mathbb{Z}_k$-tension polynomial is a divisor of the chromatic polynomial. So Stanley's Reciprocity Theorem for the chromatic polynomial gives us a reciprocity result (Theorem 3.3.6) for the $\mathbb{Z}_k$-tension polynomial.

Finding a combinatorial reciprocity theorem for the modular flow polynomial, on the other hand, was an open problem, posed in [BZ06b]. Theorem 3.3.5 solves this problem. It is our main contribution in this chapter.

3.3.5. Modular Flow Reciprocity Theorem.

Let $G = (V, E)$ be a graph with c components. Then for all $0 < k \in \mathbb{N}$

$$(-1)^{|E|-|V|+c} \cdot \bar{\varphi}_G(-k) \quad = \quad \#\{\ (f,\sigma)\ : \begin{array}{l} f \text{ a } \mathbb{Z}_k\text{-flow on } G \\ \sigma \text{ a totally cyclic orientation of } G/\mathsf{supp}(f)\ \}. \end{array}$$

We give two proofs of this theorem. The first is a geometric proof that is based on applying Ehrhart-Macdonald Reciprocity to a disjoint union of polytopes simultaneously. We give this proof in Section 3.6 and show how this methods helps finding the correct statement of the reciprocity theorem. This approach is not immediately related to inside-out polytopes; in particular we can do without a generalized version of Ehrhart-Macdonald Reciprocity for inside-out polytopes. We draw the connection to inside-out polytopes in Section 3.7. Our second proof is elementary and purely combinatorial, based on induction and the well-known deletion/contraction formula (3.2.2) for the $\mathbb{Z}_k$-flow polynomial, which we give in Section 3.8. Interestingly, the combinatorial proof is longer than the geometric one. One should bear in mind, though, that the geometric proof relies on machinery which we merely use and do not develop from first principles.

A result similar to Theorem 3.3.5 was shown independently by Babson and Beck in their as yet unpublished manuscript [BB], where they use hyperplane arrangements on the torus for the proof.

We already mentioned that, as $\chi_G(k) = k^c \bar{\theta}_G(k)$, Stanley's Reciprocity Theorem implies a reciprocity theorem for the modular tension polynomial. See also [Che07]. Knowing the reciprocity theorem for the modular flow polynomial, we can phrase the reciprocity theorem for the modular tension polynomial such that the "duality" of these two results becomes apparent. The latter arises from the former by replacing "flows" with "tensions", "totally cyclic" with "acyclic", "contraction" with "deletion" and "$|E| - (|V| - c)$" with "$|V| - c$".

3.3.6. Modular Tension Reciprocity Theorem.

Let $G = (V, E)$ be a graph with c components. Then for all $0 < k \in \mathbb{N}$

$$(-1)^{|V|-c} \cdot \bar{\theta}_G(-k) = \#\{ (t, \sigma) : t \text{ a } \mathbb{Z}_k\text{-tension on } G,\ \sigma \text{ an acyclic orientation of } G \setminus \mathsf{supp}(t) \}.$$

We now show the equivalence of the Modular Tension Reciprocity Theorem 3.3.6 and the Chromatic Reciprocity Theorem 3.3.6. In Section 3.9 we explain how our geometric proof of the Modular Flow Reciprocity Theorem can be translated into a proof of the modular tension reciprocity theorem.

Equivalence of Theorems 3.3.1 and 3.3.6. In light of the proof of 3.2.6 it suffices to argue that for a given k-coloring c and a corresponding $\mathbb{Z}_k$-tension t, the acyclic orientations of $G \setminus \mathsf{supp}(t)$ are in bijection with the acyclic orientations of G that are compatible with c. If σ is an acyclic orientation of G then clearly $\sigma|_{E\setminus\mathsf{supp}(t)}$ is an acyclic orientation of $G \setminus \mathsf{supp}(t)$. Conversely, let σ' be an acyclic orientation of $G \setminus \mathsf{supp}(t)$. We have to show that there is a unique extension of σ' to an acyclic orientation σ of G that is compatible with c. However the condition that $e = uv$ has to be oriented from u to v whenever $c(u) < c(v)$, fixes the orientation of all edges in $\mathsf{supp}(t)$. Suppose the resulting orientation σ did contain a directed cycle C. Then all vertices on C have to have the same color with respect to c, as following an edge can never decrease the color. But this means that $C \subset E \setminus \mathsf{supp}(t)$ which is a contradiction to ${}_{\sigma'}G \setminus \mathsf{supp}(t)$ being acyclic. □

3.4. *The Linear Algebra Connection*

Let us rephrase our definition of flows and tensions in terms of linear algebra. All the material in this section is well-known.

In case of $\mathbb{Z}$-flows or $\mathbb{R}$-flows, the constraint (3.1) can be written conveniently in terms of the incidence matrix A of G. The **incidence matrix** of G is the matrix $A \in \{0, \pm 1\}^{V \times E}$ with $A_{v,e} = 1$ if $\mathsf{head}(e) = v$, $A_{v,e} = -1$ if $\mathsf{tail}(e) = v$ and $A_{v,e} = 0$ otherwise. If e is a loop we define $A_{v,e} = 0$. As a consequence the incidence matrix does not determine the associated graph uniquely! If $A_{v,\cdot}$ denotes the row of the incidence matrix corresponding to v, then the instance of equation (3.1) corresponding to v can be written as $\langle A_{v,\cdot}, f \rangle = 0$. In other words, the $\mathbb{R}$-flows on G are precisely those vectors $f \in \mathbb{R}^E$ with $Af = 0$. Consequently $\ker A$ is called the **flow space** of G. The $\mathbb{Z}$-flows on G are precisely $\mathbb{Z}^E \cap \ker A$, the lattice points in the flow space.

Translates of the flow space, more precisely the affine spaces $\{f \in \mathbb{R}^E | Af = b\}$ for some integral $b \in \mathbb{Z}^V$, will play an important role in what is to come. The $b \in \text{Im } A$ have the following property.

3.4.1. If $G = (V, E)$ is a graph with c components and A its incidence matrix. Then $\dim(\text{Im } A) = |V| - c$ and

$$\text{Im } A = \{b \in \mathbb{R}^V \mid \sum_{v \in K} b_v = 0 \text{ for every component } K \text{ of } G\}.$$

In particular $\sum_{v \in V} (Af)_v = 0$ for every $f \in \mathbb{R}^E$.

Proof. For every component K of G and every column $A_{\cdot e}$ of A, $\sum_{v \in K} A_{ve} = 0$. This shows both that the rank of A is at most $|V| - c$ and that the set on the right hand side is contained in Im A. However there are only c constraints $\sum_{v \in K} b_v = 0$ and so the set on the right hand side has dimension at least $|V| - c$. □

As in the case of flows, $\mathbb{Z}$- and $\mathbb{R}$-tensions can be characterized in terms of a matrix. Given a cycle c in G and a totally cyclic orientation $\sigma : E \to \{+, -\}$ of $G[c]$, we can define a sign vector $x_c : E \to \{0, \pm 1\}$ that coincides with σ on c and is zero everywhere else, i.e. $\text{supp}(x_c) = c$ and $x_c|_c = \sigma$. Using x_c the instance of equation (3.2) corresponding to c can be written more compactly as $\langle x_c, t \rangle = 0$. Note that there are precisely two choices of sigma and the corresponding sign vectors x_c and x'_c are related by $x_c = -x'_c$ and thus $\langle x_c, t \rangle = 0$ if and only if $\langle x'_c, t \rangle = 0$. We will not make further use of the notation x_c. Instead we will identify a cycle with its sign vector (for an arbitrary fixed orientation σ) and denote both by the same symbol, e.g. c. Now let C be the $0, -1, +1$-matrix whose rows are the sign vectors of all cycles in G. Such a matrix we call a **cycle matrix**[1] of G. The columns of C are indexed by edges in E while the rows are indexed by cycles. Then the $\mathbb{R}$-tensions t of G are precisely those maps $t : E \to \mathbb{R}$ with $Ct = 0$, i.e. precisely the elements of ker C. The $\mathbb{Z}$-tensions of G are precisely the lattice points in ker C. We call ker C the **tension space** of G.

A **cut** in G is a set of edges S such that there exists a partition of the vertex set into two classes $V = X \cup Y$ such that S consists of precisely those edges incident to both a vertex in X and a vertex in Y. The sets X and Y are the **shores** of the cut. A **single vertex cut** is a cut given by a partition of the form $V = v \cup (V \setminus v)$ or $V = (V \setminus v) \cup v$. The cuts can be seen as vectors $x_S \in \{0, 1\}^E$ with $\text{supp}(x_S) = S$. The set of all cuts forms an $\mathbb{F}_2$-vector space called the **cut space** of G. We can also view a cut as a signed vector $x_S \in \{0, \pm 1\}^E$: we define x_S to be the vector with $\text{supp}(x_S) = S$ such that $x_S(e) = 1$ if e is an edge from a vertex in X to a vertex of Y and $x_S(e) = -1$ if e is an edge from a vertex in Y to a vertex in X. Such a vector we call a **signed cut** or **directed cut** and if it stems from a single vertex cut we call it a **signed single vertex cut**. However these objects we have encountered before! On the one hand, the signed cuts are precisely those tensions coming from a 2-coloring c with $c(v) = 0$ if $v \in X$ and $c(v) = 1$ if $v \in Y$. These are 2-tensions, but not all 2-tensions arise from a signed cut. The signed single vertex cuts on the other hand are precisely the rows of the incidence

[1] More precisely we call any $0, +1, -1$-matrix C whose rows span the flow space and whose kernel is the tension space a cycle matrix. We will meet a particular choice of the cycle matrix later in this section.

matrix (and their negatives). So the signed single vertex cuts are $0, +1, -1$-tensions and, by definition, any $\mathbb{R}$-flow is orthogonal to all of these.

A similar statement can be made about the sign vectors of cycles. On the one hand the sign vector σ of any cycle c is a $0, +1, -1$-flow. Let c be a cycle. Then $\sigma : E \to \{0, \pm\}$ is its sign vector if $\mathsf{supp}(\sigma) = c$ and ${}_{\sigma|_c}G[c]$ is a directed cycle. The edges in ${}_{\sigma_c}G[c]$ are oriented such that at every vertex lying on the cycle there is precisely one incoming and one outgoing edge. So (3.1) holds and thus σ is a flow with entries $0, \pm 1$. Note, though, that the unsigned characteristic vectors of cycles do not in general lie in flow space, just as unsigned characteristic vectors of cuts do not lie in tension space. On the other hand the sign vectors of cycles are just the rows of the cycle matrix. So the sign vectors of cycles are $0, +1, -1$-flows and, by definition, any $\mathbb{R}$-tension is orthogonal to all of these.

Moreover it turns out that:

> The signed single vertex cuts span the tension space and the sign vectors of cycles span the flow space.

In fact if we take all vertices except one for every component of G, then the corresponding signed single vertex cuts[2] form a basis of the tension space. And if we take the sign vectors of the cycles in any basis of the cycle space[3], we obtain a basis of the flow space. We are going to meet one family of such bases below. But first we observe that the above implies that the flow space and the tension space are orthogonal complements in edge space and we even know their dimension:

> The dimension of the tension space is $|V| - c$ and the dimension of flow space is $|E| - |V| + c$.

We conclude this section by constructing a particular basis of the flow space of a graph. Taking these vectors as rows, we obtain a particular choice of cycle matrix C that will be of use to us in the next sections.

We call a spanning forest of G that consists of one spanning tree per component of G a **spanning multi-tree**. Note that not all spanning forests of G are spanning multi-trees, hence this unusual term. Let T be a spanning multi-tree of G. For any edge e of G there is a unique path P in T such that $e \cup P = C$ is a cycle. Let σ be the totally cyclic orientation of $G[C]$ with $\sigma(e) = +$. Let $(C_e)_{E \setminus T}$ denote the collection of sign vectors of these cycles. These vectors form a basis of the flow space. Let C denote the cycle matrix that has these sign vectors as rows. We call this matrix the **cycle matrix given by the multi-tree** T.

Assume without loss of generality that the columns of C, i.e. the edges in E, are ordered such that the edges in $E \setminus T$ come first. Assume that the rows, which are indexed by $E \setminus T$ are ordered correspondingly. Then C has the form

$$C = \begin{pmatrix} I & \bar{C} \end{pmatrix}$$

where I denotes the $E \setminus T \times E \setminus T$ identity matrix. The matrix $\bar{C}$ is a matrix whose rows are indexed by $E \setminus T$ and whose columns are indexed by T. This matrix links cycle matrices to

[2] Of course for any such vertex v we only take one of the cuts $V \cup (V \setminus v)$ and $(V \setminus v) \cup V$.

[3] Again, for any cycle c we only take the sign vector for one of the two totally cyclic orientations of c.

another important class of matrices. Consider the following matrix N, with rows indexed by T and columns indexed by E.

$$N = \begin{pmatrix} -\bar{C}^t & I \end{pmatrix}$$

Here I denotes the $T \times T$ identity matrix. To interpret this definition, we observe that a multi-tree T has the property that for any edge $e_E = uv \in E$ there is a unique path P between u and v in T. Let σ be an orientation that turns ${}_\sigma G[P]$ into a directed path from u to v. If $e_E \in T$ that path is just e_E itself. If $e_E \in E \setminus T$, that path is just the reverse of the path we used to construct C. The columns of N are just the sign vectors of these paths P. We can thus describe the entries of N as follows: $N_{e_T,e_E} = +1$ if $e_T \in P$ and $\sigma(e_T) = +$, $N_{e_T,e_E} = -1$ if $e_T \in P$ and $\sigma(e_T) = -$ and $N_{e_T,e_E} = 0$ otherwise. The matrix N is known as the **network matrix** of G and the spanning multi-tree T. See [Sch86, Sec. 19.3].

We summarize the key facts stated in this section in the following theorem.

3.4.2. Theorem.

Let $G = (V, E)$ be a graph with c components. Let T be a spanning multi-tree of G. Let A be the incidence matrix of G and let C be the cycle matrix of G given by T. Then:

$$\text{tension space} = \ker C = \operatorname{Im} A^t \qquad \text{flow space} = \ker A = \operatorname{Im} C^t$$

$$\dim(\text{tension space}) = |V| - c \qquad \dim(\text{flow space}) = |E| - |V| + c$$

$$\text{tension space} = \text{flow space}^\perp$$

In the next section we are going to see that incidence matrices, network matrices and this particular choice of cycle matrix have a very important property: they are totally unimodular.

3.5. *Total Unimodularity*

In this section we introduce the concept of total unimodularity and show that the matrices we encountered in the previous section are totally unimodular. In this section we follow Schrijver's textbook [Sch86].

A matrix is **totally unimodular** if the determinant of every square submatrix is 0, +1 or −1. So in particular a totally unimodular matrix can only have entries 0, +1 or −1. This implies that a totally unimodular matrix maps lattice points to lattice points. Moreover, if $A \in \mathbb{R}^m \times n$ is a totally unimodular matrix, b is an integral vector and the affine map $x \mapsto Ax + b$ is injective, then this affine map gives a bijection between the lattice points in the domain and the lattice points in the image. This is a consequence of Cramer's Rule. So in particular, linear (or affine) automorphisms that are given by totally unimodular matrices are lattice transformations.

Why are we are we interested in total unimodularity? In the last section we associated certain matrices with our graph G. We are going to uses these matrices to construct polytopes and it will be of importance for our purposes that these polytopes are lattice polytopes. It turns out that totally unimodular matrices and lattice polytopes are closely related. If a polytope P is given by a system $Ax \leq b$ of linear inequalities where A is totally unimodular and b is integral, then P is a lattice polytope. In fact there is the following theorem.

3.5.1. Theorem. (Hoffman and Kruskal, cf. [Sch86, Theorem 19.3.iii])

Let $A \in \{0, \pm 1\}^{m \times n}$. Then A is totally unimodular if and only if for all $a, b \in \mathbb{Z}^n$ and all $c, d \in \mathbb{Z}^m$ the polytope

$$\{x \in \mathbb{R}^n | a \leq x \leq b,\ c \leq Ax \leq d\}$$

is a lattice polytope.

This theorem describes how we are going to *use* the total unimodularity of a matrix. The following is a useful tool for *showing* the total unimodularity of a matrix.

3.5.2. Theorem. (Ghouila-Houri, cf. [Sch86, Theorem 19.3.iv])

A $0, \pm 1$-matrix A is totally unimodular if and only if any collection of rows of A can be partitioned into two classes such that the sum of rows in the one class minus the sum of rows in the other class is a $0, \pm 1$-vector.

Also, there are a number of operations on matrices that preserve total unimodularity. The following lemma lists a couple of these.

3.5.3. Lemma. (cf. [Sch86, Sec. 19.4])

If A is a totally unimodular matrix and A' arises from A by permuting rows or columns, passing to the transpose, passing to a submatrix, multiplying a row or column by -1, repeating a row or column, adding an all-zero row or column or adding a row or column that contains just a single non-zero entry which is ± 1, then A' is totally unimodular.

Now for our purposes the key result is that the incidence matrix of a graph and the network and cycle matrices of a graph with respect to any spanning multi-tree are totally unimodular.

3.5.4. Theorem.

Let G be a graph and T a spanning multi-tree of G. Then:

1. The incidence matrix of G is totally unimodular.
2. The network matrix of G and T is totally unimodular.
3. The cycle matrix of G and T is totally unimodular.

Proof. Both 1. and 2. can be found in [Sch86, Sec. 19.4]. 1. is a direct consequence of the characterization of Ghouila-Houri (Theorem 3.5.2) while 2. requires a bit more work. 3. can be seen as follows, using the notation from the end of Section 3.4 and applying Lemma 3.5.3 repeatedly. As the network matrix N of G and T is totally unimodular, so is its submatrix $-\bar{C}^t$. Taking the transpose and multiplying every row by -1 we see that $\bar{C}$ is totally unimodular. As adding columns that contain just is single 1 preserves total unimodularity, we obtain that C is totally unimodular. □

Note that arbitrary cycle matrices, that is matrices whose rows correspond to sign vectors of cycles that span the flow space, are not totally unimodular. Theorem 3.5.4 states that those cycle matrices coming from a spanning multi-tree are.

3.6. A Geometric Proof of Modular Flow Reciprocity

In this section we give a geometric proof of the Modular Flow Reciprocity Theorem 3.3.5 using the Ehrhart-Macdonald Reciprocity Theorem 3.3.2.

Before we begin, let us deal with the pathological cases first. Suppose that G has a bridge e. Then there are no nowhere zero $\mathbb{Z}_k$-flows on G by 3.2.4 and $\bar{\varphi}(k)_G = 0$. On the other hand, any $\mathbb{Z}_k$-flow f on G has $e \in \mathsf{zero}(f)$ and so $G/\mathsf{supp}(f)$ still contains a bridge. By 3.2.3 this implies that $G/\mathsf{supp}(f)$ has no totally cyclic orientation. Thus Theorem 3.3.5 holds if G has a bridge. So without loss of generality we can assume that G does not have a bridge and begin with the proof proper.

The first step is to model the set of all nowhere zero $\mathbb{Z}_k$-flows of a graph $G = (V, E)$ geometrically. To that end we want to view a $\mathbb{Z}_k$-flow $f : E \to \mathbb{Z}_k$ as a map $f : E \to \mathbb{Z}$, i.e. as a lattice point in $\mathbb{Z}^E$.

Two maps $f_{\mathbb{Z}} : E \to \mathbb{Z}$ and $f_{\mathbb{Z}_k} : E \to \mathbb{Z}_k$ are said to **correspond** if $f_{\mathbb{Z}_k}(e)$ is the coset of $f_{\mathbb{Z}}(e)$ in $\mathbb{Z}_k = \mathbb{Z}_{/k\mathbb{Z}}$ for all $e \in E$. For every $f_{\mathbb{Z}}$ there is a unique corresponding $f_{\mathbb{Z}_k}$. For every $f_{\mathbb{Z}_k}$ there is a corresponding $f_{\mathbb{Z}}$, but it is not uniquely determined. $f_{\mathbb{Z}}$ is uniquely determined if we require $0 \le f_{\mathbb{Z}} < k$, i.e. if we identify the cosets in $\mathbb{Z}_k$ with their respective representatives $0, \ldots, k-1 \in \mathbb{Z}$. Unless otherwise stated, this is the map we mean by "the corresponding $f_{\mathbb{Z}}$". Note that this gives a bijection between the nowhere zero maps $f_{\mathbb{Z}_k} : E \to \mathbb{Z}_k$ and the vectors $f_{\mathbb{Z}} \in \mathbb{Z}^E$ with $0 < f_{\mathbb{Z}} < k$, that is, the lattice points in the open cube $(0, k)^E$.

A $\mathbb{Z}_k$-flow is a map $f_{\mathbb{Z}_k} : E \to \mathbb{Z}_k$, such that for every vertex $v \in V$ equation (3.1) holds. Let A denote the incidence matrix of G and $f_{\mathbb{Z}} : E \to \mathbb{Z}$ a map corresponding to $f_{\mathbb{Z}_k}$. Then

$$
\begin{array}{llllll}
 & \forall v \in V: & & \sum_{e=uv\in E} f_{\mathbb{Z}_k}(e) - \sum_{e=vu\in E} f_{\mathbb{Z}_k}(e) & = 0 & \text{in } \mathbb{Z}_k \\
\Leftrightarrow & \forall v \in V: & & \sum_{e=uv\in E} f_{\mathbb{Z}}(e) - \sum_{e=vu\in E} f_{\mathbb{Z}}(e) & \equiv 0 & \text{mod } k \\
\Leftrightarrow & \forall v \in V: & \exists b_v \in \mathbb{Z}: & \sum_{e=uv\in E} f_{\mathbb{Z}}(e) - \sum_{e=vu\in E} f_{\mathbb{Z}}(e) & = b_v k & \text{in } \mathbb{Z} \\
\Leftrightarrow & & \exists b \in \mathbb{Z}^V: & A f_{\mathbb{Z}} & = kb &
\end{array}
$$

and so:

3.6.1. $f_{\mathbb{Z}} : E \to \mathbb{Z}$ has $Af_{\mathbb{Z}} = kb$ for some $b \in \mathbb{Z}^V$ if and only if the corresponding $f_{\mathbb{Z}_k} : E \to \mathbb{Z}_k$ is a $\mathbb{Z}_k$-flow.

In particular we can identify the nowhere zero $\mathbb{Z}_k$-flows on G with the lattice points $f \in \mathbb{Z}^E$ with $0 < f < k$ for which there exists a $b \in \mathbb{Z}^V$ with $Af = kb$. If we now define for any $b \in \mathbb{Z}^V$ the open polytope P_b° by

$$P_b^\circ \quad := \quad \left\{ f \in \mathbb{R}^E \;\middle|\; 0 < f < 1, Af = b \right\}$$

then the set of nowhere zero $\mathbb{Z}_k$-flows is in bijection with the set

$$\bigcup_{b\in\mathbb{Z}^V} \left\{ f \in \mathbb{Z}^E \;\middle|\; 0 < f < k, Af = kb \right\} \quad = \quad \bigcup_{b\in\mathbb{Z}^V} \mathbb{Z}^E \cap kP_b^\circ. \tag{3.3}$$

The set P_b° is the intersection of the open unit cube $(0,1)^E$ with the affine subspace $\{f|Af = b\}$, which is an integral translate of the flow space $\{f|Af = 0\}$. Now, as the affine subspaces $\{f|Af = b\}$ for $b \in \mathbb{Z}^V$ are all parallel, the open polytopes P_b° are pairwise disjoint. Moreover, as $|(Af)_v| \leq \deg(v)$ for any f with $0 \leq f \leq 1$ and any vertex $v \in V$, only finitely many of these subspaces actually meet the cube, that is, only finitely many of these open polytopes are non-empty. This means in particular that the union on the right-hand side of (3.3) is disjoint and finite. We call those $b \in \mathbb{Z}^V$ such that $P_b^\circ \neq \emptyset$ **feasible** and denote the set of all feasible b by $\mathcal{B}_G$. Using this notation and (3.3) we have

3.6.2.

$$\bar{\varphi}(k) = \# \bigcup_{b\in\mathcal{B}_G} \mathbb{Z}^E \cap kP_b^\circ = \sum_{b\in\mathcal{B}_G} \mathsf{L}_{P_b^\circ}(k).$$

Now we wish to apply Ehrhart-Macdonald Reciprocity to each of the Ehrhart functions $\mathsf{L}_{P_b^\circ}(k)$ for $b \in \mathcal{B}_G$. To this end, let us denote the intersection of the same affine subspace with the closed unit cube $[0,1]^E$ by

$$P_b \quad := \quad \left\{ f \in \mathbb{R}^E \;\middle|\; 0 \leq f \leq 1, Af = b \right\}. \tag{3.4}$$

Note that these are closed polytopes and, by the same argument as above, they are pairwise disjoint.

Before we continue with the proof, we take a look at two examples.

Example A. Let G be the following graph:

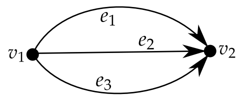

The edge space of G is three dimensional. The dimension of the flow space of G is $|E| - |V| + c = 3 - 2 + 1 = 2$. The incidence matrix of G is

$$A_1 = \begin{pmatrix} -1 & -1 & -1 \\ 1 & 1 & 1 \end{pmatrix}.$$

So $P_b^\circ = \{x \in \mathbb{R}^3 : 0 < x_1, x_2, x_3 < 1, -x_1 - x_2 - x_3 = b_1, x_1 + x_2 + x_3 = b_2\}$, which is non-empty iff $b \in \{(-1,1), (-2,2)\} = \mathcal{B}_G$. The following figure shows $P_{(-1,1)}$ and $P_{(-2,2)}$ as slices of the cube $[0,1]^3$.

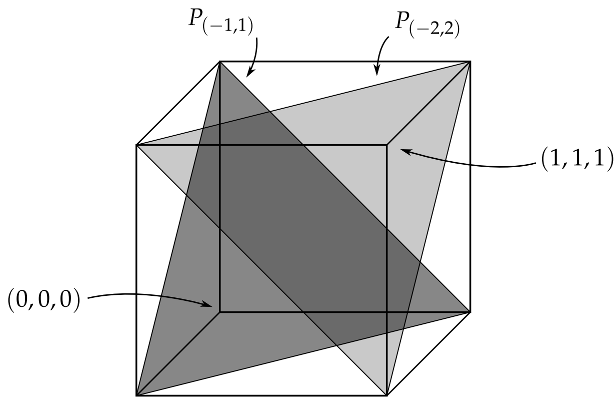

Note that the hyperplanes $\{x \in \mathbb{R}^3 : Ax = (0,0)\}$ and $\{x \in \mathbb{R}^3 : Ax = (3,3)\}$ also intersect the cube $[0,1]^3$ in the points $(0,0,0)$ and $(1,1,1)$, respectively, but they do not meet the open cube $(0,1)^3$.

Example B Consider the following graph on 3 vertices.

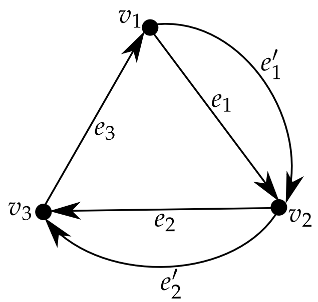

The edge space of G is 5 dimensional and its incidence matrix is

$$A = \begin{pmatrix} -1 & -1 & 0 & 0 & 1 \\ 1 & 1 & -1 & -1 & 0 \\ 0 & 0 & 1 & 1 & -1 \end{pmatrix}.$$

where the columns are indexed with the edges in the order $e_1, e'_1, e_2, e'_2, e_3$. The flow space of G has dimension $|E| - |V| + c = 5 - 3 + 1 = 3 = \dim\ker A$. We can identify the set of nowhere zero $\mathbb{Z}_k$-flows with the set of lattice points in $\bigcup_{b \in \mathcal{B}_G} kP_b^\circ$. Again, the question arises what $\mathcal{B}_G$ is, i.e. for which $b \in \mathbb{Z}^V$ the affine space $\{f : A_G f = b\}$ intersects the open cube $(0,1)^E$. As $0 < f < 1$ the in- and out-degrees of v_1 and v_2 imply $-1 \le b_{v_1} \le 0$ and $0 \le b_{v_2} \le 1$, and indeed

$$\mathcal{B}_G = \left\{ \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ -1 \\ 1 \end{pmatrix}, \begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix}, \begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix} \right\}$$

as witnessed, respectively, by the vectors

$$\begin{pmatrix} 1/4 \\ 1/4 \\ 1/4 \\ 1/4 \\ 1/2 \end{pmatrix}, \begin{pmatrix} 1/4 \\ 1/4 \\ 3/4 \\ 3/4 \\ 1/2 \end{pmatrix}, \begin{pmatrix} 3/4 \\ 3/4 \\ 1/4 \\ 1/4 \\ 1/2 \end{pmatrix}, \begin{pmatrix} 3/4 \\ 3/4 \\ 3/4 \\ 3/4 \\ 1/2 \end{pmatrix}.$$

These correspond to the following four flows $f : E \to \mathbb{R}$ on G.

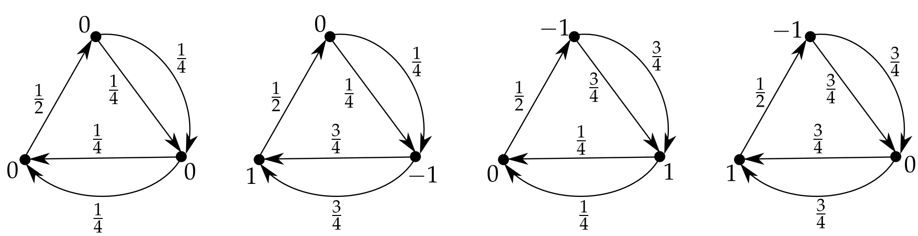

As $\bigcup_{b\in\mathcal{B}_G} kP_b^\circ$ is a union of four 3-dimensional polytopes in $\mathbb{R}^5$, we cannot draw a picture just now. However, the inside-out polytope constructed in Section 3.7, which is a projection of $\bigcup_{b\in\mathcal{B}_G} kP_b^\circ$, will allow us to visualize this set.

Now back to the proof. To be able to apply Ehrhart-Macdonald Reciprocity we need to gather more information about our polytopes.

First, we need to compute the dimension of the P_b°. This is straightforward. The set P_b° is open with respect to the affine subspace $\{f|Af = b\}$. So either P_b° is empty, or it has the same dimension as $\{f|Af = b\}$. This affine subspace is a translate of the flow space, which has dimension $|E| - |V| + c$.

3.6.3. If $b \in \mathcal{B}_G$, then $\dim P_b^\circ = |E| - |V| + c$.

So all the open polytopes P_b° that are non-empty have the same dimension. This is not true of the closed polytopes P_b. There are integral translates of flow space that intersect the unit cube only in its boundary and give rise to non-empty polytopes P_b of lower dimension, in which case the corresponding P_b° are empty.

Next, we need to check that P_b° is indeed the relative interior of P_b. By our observations about the dimension, this does not hold for arbitrary $b \in \mathbb{Z}^V$, but if $b \in \mathcal{B}_G$, this is true. If $P_b^\circ \neq \emptyset$, then $\mathrm{aff}(P_b^\circ) = \mathrm{aff}(P_b)$ and as P_b° is relatively open, $P_b^\circ \subset \mathrm{relint}\, P_b$. On the other hand let $f \in P_b$ with $f(e) \in \{0, k\}$ for some e. Let $\epsilon > 0$. Let z be a nowhere zero $\mathbb{R}$-flow on G with $||z|| < \epsilon$. Then $f + z, f - z \in \mathrm{aff}(P_b)$ but one of the two is not in P_b. So f is not in the relative interior of P_b which proves $P_b^\circ = \mathrm{relint}\, P_b$.

3.6.4. If $b \in \mathcal{B}_G$, then $P_b^\circ = \mathrm{relint}\, P_b$.

Finally, we need to verify that the P_b are lattice polytopes.[4] But for this we can use the machinery we presented in Section 3.5. By Theorem 3.5.4 the incidence matrix A is totally unimodular and by the Theorem of Hoffman and Kruskal 3.5.1:

3.6.5. If $b \in \mathbb{Z}^V$, the polytopes P_b are lattice polytopes.

Now that we have all the ingredients together we can apply Ehrhart-Macdonald Reciprocity 3.3.2 and compute

$$\begin{aligned}\bar{\varphi}(-k) &= \sum_{b\in\mathcal{B}_G} \mathsf{L}_{P_b^\circ}(-k) = \sum_{b\in\mathcal{B}_G} (-1)^{|E|-|V|+c} \mathsf{L}_{P_b}(k) = (-1)^{|E|-|V|+c} \cdot \sum_{b\in\mathcal{B}_G} \mathsf{L}_{P_b}(k) \\ &= (-1)^{|E|-|V|+c} \cdot \# \bigcup_{b\in\mathcal{B}_G} \mathbb{Z}^E \cap kP_b.\end{aligned}$$

So what Ehrhart-Macdonald Reciprocity tells us, is that if we evaluate $\bar{\varphi}$ at a negative value, instead of counting the lattice points in a disjoint union of open polytopes, we now have to count the lattice points in the disjoint union of their closures.

[4] Note that this step is not strictly necessary. Ehrhart-Macdonald Reciprocity also holds for rational polytopes in which case the Ehrhart function is a quasi-polynomial. However, it will be of use in Chapter 4 to know that the P_b are lattice polytopes. So we do some extra work here and use the simpler version of Ehrhart-Macdonald Reciprocity for lattice polytopes.

3.6.6.
$$(-1)^{|E|-|V|+c} \cdot \bar{\varphi}(-k) = \# \bigcup_{b \in \mathcal{B}_G} \mathbb{Z}^E \cap kP_b$$

All we have left to do is give an interpretation for the lattice points in this union, i.e. in the set on the right hand side of 3.6.6.

Loosely speaking, we would like to identify a point $f \in \mathbb{Z}^E \cap kP_b$ with the corresponding $\mathbb{Z}_k$-flow $f_{\mathbb{Z}_k}$. But this correspondence is not one-to-one anymore, as $f(e)$ can take on either of the values 0 and k, which both represent the neutral element of $\mathbb{Z}_k$. So in a way, the points in $\mathbb{Z}^E \cap kP_b$ represent $\mathbb{Z}_k$-flows with "two kinds of zeros". We are going to model these objects as pairs of $\mathbb{Z}_k$-flows $f_{\mathbb{Z}_k}$ and certain sign patterns $t : \mathsf{zero}(f_{\mathbb{Z}_k}) \to \{+,-\}$. The question is, which sign patterns appear.

To implement these ideas more formally, we start off with a different question. Let $f \in \mathbb{Z}^E$ with $0 \le f \le k$. Then automatically $Af \in \mathbb{Z}^V$. How can we tell whether $Af =: b$ is feasible, that is, whether $f \in P_b$ for $b \in \mathcal{B}_G$? b is feasible if and only if there exists an $f' \in P_b^\circ$. For such an f' the vector $z := f' - f$ is an $\mathbb{R}$-flow by $Az = Af' - Af = b - b = 0$. Moreover, as $0 < f' < k$, if $f(e) = 0$ then $z(e) > 0$ and if $f(e) = k$ then $z(e) < 0$. So f determines the sign pattern of z on $\{e | f(e) \equiv 0 \mod k\}$. Conversely, if z is an $\mathbb{R}$-flow with $f(e) = 0 \Rightarrow z(e) > 0$ and $f(e) = k \Rightarrow z(e) < 0$, then there exists some small $\delta > 0$ such that $f + \delta z \in P_b^\circ$. We summarize:

3.6.7. Let $f \in \mathbb{Z}^E$ with $0 \le f \le k$. $f \in P_b$ for $b \in \mathcal{B}_G$ if and only if there exists an $\mathbb{R}$-flow z on G such that $f(e) = 0 \Rightarrow z(e) > 0$ and $f(e) = k \Rightarrow z(e) < 0$.

It appears that the sign pattern of z on the edges $\{e | f(e) \not\equiv 0 \mod k\}$ is irrelevant. This can be made precise in the following way.

Let $S \subset E$. If z is a Grp-flow on G, then $z|_{E \setminus S}$ is a Grp flow on G/S. Conversely, if z' is a Grp-flow on G/S, then there exists a Grp-flow z on G with $z|_{E \setminus S} = z'$, cf. Lemmas 3.8.2 and 3.8.3 in Section 3.8. So an $\mathbb{R}$-flow z as given in 3.6.7 implies the existence of an $\mathbb{R}$-flow z' on $G/\{e | f(e) \not\equiv 0 \mod k\}$ such that $f(e) = 0 \Rightarrow z(e) > 0$ and $f(e) = k \Rightarrow z(e) < 0$. Conversely such an $\mathbb{R}$-flow z' implies the existence of an $\mathbb{R}$-flow z on G as given in 3.6.7. Thus, we come to the following observation.

3.6.8. Let $f \in \mathbb{Z}^E$ with $0 \le f \le k$. $f \in P_b$ for some $b \in \mathcal{B}_G$ if and only if there exists a nowhere zero $\mathbb{R}$-flow z on $G/\{e | f(e) \not\equiv 0 \mod k\}$ such that $f(e) = 0 \Rightarrow z(e) > 0$ and $f(e) = k \Rightarrow z(e) < 0$.

We can now complete the proof of the Modular Flow Reciprocity Theorem 3.3.5.

Proof of 3.3.5.. We define a map

$$h : \bigcup_{b \in \mathcal{B}_G} \mathbb{Z}^E \cap kP_b \quad \to \quad \#\{\, (f, \sigma) : \begin{array}{l} f \text{ a } \mathbb{Z}_k\text{-flow on } G \\ \sigma \text{ a totally cyclic orientation of } G/\mathsf{supp}(f) \end{array} \}$$

by $h(f_{\mathbb{Z}}) = (f_{\mathbb{Z}_k}, \sigma)$ where $f_{\mathbb{Z}_k}$ is the $\mathbb{Z}_k$-flow on G corresponding to $f_{\mathbb{Z}}$ and $\sigma : \mathsf{zero}(f) \to E$ is given by $\sigma(e) = +$ if $f'(e) = 0$ and $\sigma(e) = -$ if $f'(e) = k$. We are going to show that h is a bijection, which completes the proof by 3.6.6.

h is well-defined. By 3.6.1 if $f_{\mathbb{Z}} \in kP_b$ for any $b \in \mathbb{Z}^V$, then $f_{\mathbb{Z}_k}$ is a $\mathbb{Z}_k$-flow. By 3.6.8, σ is the sign pattern of a nowhere zero $\mathbb{R}$-flow on $G/\mathsf{supp}(f)$. So by 3.2.3, σ is a totally cyclic orientation of $G/\mathsf{supp}(f)$.

h is injective. Every pair $(f_{\mathbb{Z}_k}, \sigma)$ has a unique preimage $f_{\mathbb{Z}} \in \mathbb{Z}^E$ with $0 \leq f_{\mathbb{Z}} \leq k$ under h by construction. $f_{\mathbb{Z}_k}$ determines the values of $f_{\mathbb{Z}}$ on $\mathsf{supp}(f_{\mathbb{Z}_k})$ and σ determines the values of $f_{\mathbb{Z}}$ on $\mathsf{zero}(f_{\mathbb{Z}_k})$.

h is surjective. Let $(f_{\mathbb{Z}_k}, \sigma)$ be a pair of a $\mathbb{Z}_k$-flow $f_{\mathbb{Z}_k}$ and a totally cyclic orientation σ of $G/\mathsf{supp}(f)$. Let $f_{\mathbb{Z}}$ denote the unique map corresponding to $f_{\mathbb{Z}_k}$ with $0 \leq f_{\mathbb{Z}} \leq k$ and $f_{\mathbb{Z}}(e) = 0$ if $\sigma(e) = +$ and $f_{\mathbb{Z}}(e) = k$ if $\sigma(e) = -$. By 3.6.1, $Af = kb$ for some $b \in \mathbb{Z}^V$. As σ is totally cyclic, there exists, by 3.2.3, a nowhere zero $\mathbb{R}$-flow z on $G/\mathsf{supp}(f)$ such that $\mathsf{sgn}(z) = \sigma$. So by 3.6.8, $b \in \mathcal{B}_G$. □

3.7. *The Connection to Inside-Out Polytopes*

An inside-out polytope is the set of points inside a polytope and outside a hyperplane arrangement. In this section we connect our geometric proof of Modular Flow Reciprocity to the concept of inside-out polytopes. The benefit will be a geometric explanation of the fact that the leading coefficient of the modular flow polynomial of a graph G is 1.

An **inside-out polytope** is a pair $(P, \mathcal{H})$ of an open rational d-dimensional polytope P and a finite transverse hyperplane arrangement $\mathcal{H}$ in $\mathbb{R}^n$. $\mathcal{H}$ is **transverse** to P if every hyperplane in $\mathcal{H}$ meets the relative interior of P. The connected components of $P \setminus \bigcup \mathcal{H}$ are open d-dimensional polytopes, which we call the **cells** of the inside-out polytope. The respective closures of the cells we call the **closed cells**. The (open) Ehrhart function of the inside-out polytope $(P, \mathcal{H})$ is $\mathsf{L}_{P \setminus \bigcup \mathcal{H}}(k) = \#\mathbb{Z}^n \cap k(P \setminus \bigcup \mathcal{H})$. If the closed cells of $(P, \mathcal{H})$ are are lattice polytopes, then $\mathsf{L}_{P \setminus \bigcup \mathcal{H}}(k)$ is a polynomial.

The notion of an inside-out polytope was introduced by Beck and Zaslavsky. For more information we recommend the articles [BZ06a, BZ06b]. One of their main results was the extension of Ehrhart-Macdonald Reciprocity to Ehrhart polynomials of inside-out polynomials [BZ06a, Theorem 4.1]: $(-1)^{\dim P} \mathsf{L}_{P \setminus \bigcup \mathcal{H}}(-k)$ counts the number of lattice points in $k\bar{P}$ where $\bar{P}$ is the closure of P and any point $z \in \mathbb{Z}^n \cap k\bar{P}$ is counted with multiplicity equal to the number of closed cells that contain z. This follows directly from applying Ehrhart-Macdonald reciprocity to each cell of $(P, \mathcal{H})$ simultaneously, just as we did in our geometric proof in Section 3.6. The difference is that the "closed cells" in our case were disjoint and so we did not have to count lattice points with a multiplicity. However, the question arises whether we can can construct an inside-out polytope whose Ehrhart function is the modular flow polynomial of G.

The chromatic polynomial of a graph was captured as the Ehrhart function of an inside out polynomial by Beck and Zaslavsky in [BZ06a], who achieved the same thing for the integral flow polynomial in [BZ06b]. The case of the integral tension polynomial was settled by Dall in [Dal08]. In this section we describe an inside-out polytope whose Ehrhart function is the modular flow polynomial of G and the same can be done for the modular tension polynomial, see Section 3.9. Independently from us, similar work on the modular flow polynomial is being done by Babson and Beck [BB].

Apart from the interest of relating different points of view, there is another motivation for connecting our geometric proof to inside-out Ehrhart theory. We wish to give a geometric explanation of the following well-known fact.

3.7.1. For any graph $G = (V, E)$ with c components, the degree of $\bar{\varphi}_G$ is $|E| - |V| + c$ and the leading coefficient of $\bar{\varphi}_G$ is 1.

That $\deg(\bar{\varphi}_G) = |E| - |V| + c$ is immediate from our modeling of $\bar{\varphi}_G$ as an Ehrhart function, for the degree of L_P for a polytope P is $\dim P$ and we already noted that the dimension of the P_b° is $|E| - |V| + c$. That the leading coefficient is 1, however, is not clear yet from the geometric point of view. Combinatorially, it is easy to see from the deletion/contraction formula 3.2.2: when computing $\bar{\varphi}_G$ recursively, the only case in 3.2.2 that may potentially affect the leading coefficient is the one where e is neither a loop nor a bridge and $\bar{\varphi}_G(k) = \bar{\varphi}_{G/e}(k) - \bar{\varphi}_{G\setminus e}(k)$. However, the degree of the former polynomial is $|E| - 1 - |V| + 1 + c = |E| - |V| + c$, while the degree of the latter polynomial is $|E| - 1 - |V| + c < |E| - |V| + c$. So the leading coefficient of $\bar{\varphi}_G$ is the leading coefficient of $\bar{\varphi}_{G/e}$ which, by induction, is 1.

Let us now begin constructing an inside-out polytope for a given graph G. Let T be a spanning multi-tree of G. Let C denote the cycle matrix of G and T. We already mentioned that the flow space of G is generated by the rows of C, that is, Im C^t is the flow space of G. Recall that the columns of C corresponding to the non-tree edges contain just a single 1. So the image of the open cube $(0,k)^{E\setminus T}$ under C^t is precisely the set of those $\mathbb{R}$-flows f with $0 < f(e) < k$ for all non-tree edges $e \in E \setminus T$. A point $z \in (0,k)^{E\setminus T}$ is mapped to an $\mathbb{R}$-flow $C^t z$ with $C^t z(e) \not\equiv 0 \mod k$ for any $e \in E$ if and only if $C^t z \neq bk$ for any $b \in \mathbb{Z}^E$. So let $\mathcal{H}$ denote the arrangement of all hyperplanes $\{z | \langle C_{\cdot e}, z\rangle = bk\}$ for $b \in \mathbb{Z}^E$ and $e \in E \setminus T$ that meet the cube $(0,k)^{E\setminus T}$.

3.7.2. C^t gives a bijection between the set $\mathbb{Z}^{E\setminus T} \cap k((0,k)^{E\setminus T} \setminus \bigcup \mathcal{H})$ of lattice points in the k-th dilate of the inside-out polytope $((0,1)^{E\setminus T}, \mathcal{H})$ and the set of $\mathbb{Z}$-flows f with $0 < f(e) < k$ for all non-tree edges $e \in E \setminus T$ such that $f(e) \not\equiv 0 \mod k$ for all edges $e \in E$.

We now claim that these $\mathbb{Z}$-flows are in bijection with $\bigcup_{\mathcal{B}_G} \mathbb{Z}^E \cap kP_b^\circ$ and thus with the nowhere zero $\mathbb{Z}_k$-flows on G.

Now consider the map $r_k : \mathbb{R}^E \to \mathbb{R}^E$ that sends an $f \in \mathbb{R}^E$ to the corresponding map $r_k(f)$ with $r_k(f)(e) \equiv f(e) \mod k$ and $0 \le r_k(f) < k$. If f is an $\mathbb{R}$-flow, then $r_k(f) \in \bigcup_{b\in\mathbb{Z}^V} kP_b$ because $Af(e) \equiv Ar_k(f)(e) \mod k$. And for the same reason if f is an $\mathbb{R}$-flow with $Af(e) \not\equiv 0 \mod k$ for all $e \in E$, then $r_k(f) \in \bigcup_{b\in\mathbb{Z}^V} kP_b^\circ$. Also, r_k maps lattice points to lattice points.

Restricted to the set of $\mathbb{R}$-flows with $0 \leq f(e) < k$ for all non-tree edges $e \in E \setminus T$ the map r_k is injective, by the observation that fixing the values of those edges determines the $\mathbb{R}$-flow uniquely.

We now argue that, when restricted to the set of $\mathbb{Z}$-flows f with $0 < f(e) < k$ for all non-tree edges $e \in E \setminus T$ such that $f(e) \not\equiv 0 \mod k$ for all edges $e \in E$, r_k is surjective onto the set $\bigcup_{b\in\mathcal{B}_G} \mathbb{Z}^E \cap kP_b^\circ$. Let $f' \in \bigcup_{b\in\mathcal{B}_G} \mathbb{Z}^E \cap kP_b^\circ$. Now we construct another $\mathbb{Z}$-flow f_1 with $r_k(f_1) = r_k(f')$ such that

1. Af_1 has strictly fewer non-zero entries than Af',
2. $0 < f(e) < k$ for all non-tree edges $e \in E \setminus T$ and
3. $f(e) \not\equiv 0 \mod k$ for all edges $e \in E$.

Then the claim follows by induction. Let K be any component of G. Let $v \in K$ be a vertex such that $|Af'(v)|$ is minimal nonzero within K. Without loss of generality suppose $Af'(v) > 0$. By 3.4.1 there exists a w in K such that $Af'(w) < 0$. Then there is a path P in T from v to w. Let σ be an orientation of P such that ${}_\sigma P$ is a directed path from v to w. Define $z : E \to \mathbb{R}$ by $z(e) = Af'(v)$ if $\sigma(e) = +$ and $z(e) = -Af'(v)$ if $\sigma(e) = -$ and $z(e) = 0$ otherwise. Put $f_1 := f' + z$. Now $Af_1(v) = 0$, $Af_1(w)$ is still a multiple of k as $Af'(v)$ is a multiple of k and $Af_1(u) = Af'(u)$ for all other vertices u. So f_1 is a $\mathbb{Z}$-flow, Af_1 has strictly fewer non-zero entries than Af', and $f'(e) \equiv f_1(e) \mod k$ for all $e \in E$. Finally $0 < f_1(e) < k$ for all non-tree edges $e \in E \setminus T$ because f_1 and f' differ only on tree edges $e \in T$.

We have now proved the following statement.

3.7.3. r_k gives a bijection between the $\mathbb{Z}$-flows f with $0 < f(e) < k$ for all non-tree edges $e \in E \setminus T$ such that $f(e) \not\equiv 0 \mod k$ for all edges $e \in E$ and $\bigcup_{b\in\mathcal{B}_G} \mathbb{Z}^E \cap kP_b^\circ$.

Combining 3.7.2 and 3.7.3 with 3.6.2 we obtain that the Ehrhart function of the inside out polytope $((0,1)^{E\setminus T}, \mathcal{H})$ is just the modular flow polynomial.

3.7.4. $\bar{\varphi}_G(k) = \mathsf{L}_{(0,1)^{E\setminus T}\setminus\bigcup\mathcal{H}}(k)$

Applying the inside-out version of Ehrhart-Macdonald Reciprocity leads to the Modular Flow Reciprocity Theorem. What interests us now, however, is the interpretation of the leading coefficient of $\bar{\varphi}_G$. Here the crucial ingredient is that the leading coefficient of the Ehrhart function L_P is just the volume of the polytope P and this result carries over to inside-out polytopes, see [BR07, Corollary 3.20],[BZ06a, Theorem 4.1]. Thus the leading coefficient of $\bar{\varphi}$ is just the volume of $(0,1)^{E\setminus T}$ as a subset of $\mathbb{R}^{E\setminus T}$ which is 1.

Note also that the above construction gives a bijection between the open polytopes P_b° for $b \in \mathcal{B}_G$ and the cells of our inside-out polytope. Moreover if we view $\mathbb{R}^{E\setminus T}$ as a linear subspace of $\mathbb{R}^E$ then any cell of the inside-out polytope is lattice equivalent to the corresponding P_b°.

Let us reconsider our two examples from this point of view.

Example A continued. First, we have to pick a spanning tree of G. We choose the spanning tree $T = e_3$, so $E \setminus T = e_1, e_2$. So the corresponding cycle matrix C has two rows, indexed by e_1 and e_2, and three columns, indexed by E. The row $C_{e_1\cdot}$ corresponds to the cycle e_1, e_3 in G. In the totally cyclic orientation σ of this cycle with $\sigma(e_1) = +1$ we have $\sigma(e_3) = -1$. Similarly, the row $C_{e_2\cdot}$ corresponds to the cycle e_2, e_3 and again in a totally cyclic orientation σ of this cycle with $\sigma(e_2) = +1$ we have $\sigma(e_3) = -1$. Thus

$$C = \begin{pmatrix} 1 & 0 & -1 \\ 0 & 1 & -1 \end{pmatrix}.$$

The nowhere zero $\mathbb{Z}_k$-flows are in bijection with the lattice points inside the k-th dilate of the open unit cube where the hyperplanes $\{z : \langle C_{\cdot e}, z\rangle = b\}$ have been removed for all $e \in T$ and all $b \in \mathbb{Z}$. In other words the modular flow polynomial is the Ehrhart function of this open inside-out polytope. How does the inside-out polytope look in this example? As $|E \setminus T| = 2$, we are dealing with a 2-dimensional open unit cube. As $T = \{e_3\}$, we only have to consider one column $C_{\cdot e_3}$ and so all we have to remove are hyperplanes $\{z : -z_{e_1} - z_{e_2} = b\}$ for $b \in \mathbb{Z}$. However these hyperplanes only intersect the open cube if $b = -1$. So our inside-out polytope is the open square from which the diagonal given by $z_{e_1} + z_{e_2} = 1$ has been removed.

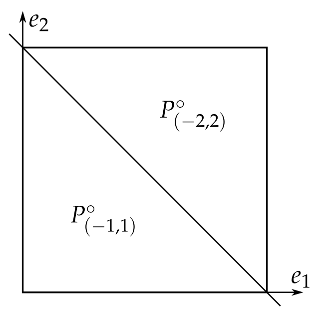

Notice how this inside-out polytope is just a projection of our disjoint union $P^\circ_{-1,1} \cup P^\circ_{-2,2}$ onto the coordinates e_1, e_2!

Example B continued. In this example we pick the spanning tree to be given by edges e_2' and e_3. The cycles given by the non tree edges e_1, e_1', e_2 are then, respectively, $e_1 e_2' e_3$, $e_1' e_2' e_3$ and $e_2 e_2'$. The orientations of the first two cycles are already totally cyclic, while in the totally cyclic orientation σ of $e_2 e_2'$ with $\sigma(e_2) = +1$ we have $\sigma(e_2') = -1$. This gives rise to the cycle matrix

$$C = \begin{pmatrix} 1 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & -1 & 0 \end{pmatrix}$$

where, again, the columns are indexed in the order $e_1, e_1', e_2, e_2', e_3$. So our inside out polytope is the 3-dimensional open unit cube where the hyperplanes given by $\langle C_{\cdot e_3}, z\rangle = z_{e_1} + z_{e'_1} = b_{e_3}$ and $\left\langle C_{\cdot e'_2}, z\right\rangle = z_{e_1} + z_{e'_1} - z_{e_2} = b_{e'_2}$ have been removed. The first family of hyperplanes intersects the cube only for $b_{e_3} = 1$, while the second family of hyperplanes intersects the cube for $b_{e'_2} \in \{0,1\}$. The intersection of the hyperplane $\langle C_{\cdot e_3}, z\rangle = 1$ with the open cube is the open quadrangle with vertices $\{(1,0,0),(0,1,0),(1,0,1),(0,1,1)\}$. The two hyperplanes $\left\langle C_{\cdot e'_2}, z\right\rangle \in \{0,1\}$ intersect the open cube in two open triangles with vertex sets $\{(0,0,0),(1,0,1),(0,1,1)\}$ and $\{(1,0,0),(0,1,0),(1,1,1)\}$, respectively.

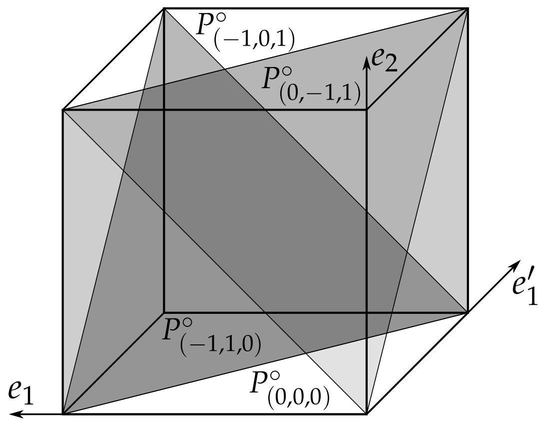

The four full dimensional cells of this inside-out polytope are projections of the four open polytopes $P^\circ_{0,0,0}, P^\circ_{0,-1,1}, P^\circ_{-1,1,0}, P^\circ_{-1,0,1}$ onto the coordinates e_1, e'_1, e_2.

3.8. *An Inductive Proof of Modular Flow Reciprocity*

In this section we give a combinatorial proof of Theorem 3.3.5. Our approach is straightforward: we show that the reciprocal of the modular flow polynomial fulfills the correct deletion/contraction formula.

Let $\mathsf{F}_G(k)$ denote the set of all pairs (f,σ) of a $\mathbb{Z}_k$-flow f and a totally cyclic orientation σ of $G/\mathsf{supp}(f)$ and let $\bar{\varphi}^*_G(k) := \#\mathsf{F}_G(k)$. Using this notation Theorem 3.3.5 simply states

$$(-1)^{\xi(G)}\bar{\varphi}_G(-k) = \bar{\varphi}^*_G(k). \tag{3.5}$$

We call $\bar{\varphi}^*_G$ the reciprocal of the modular flow polynomial. In light of 3.2.2 it suffices to show that the deletion/contraction formula given in the following theorem holds for $\bar{\varphi}^*$.

3.8.1. Theorem.

Let $G = (V, E)$ be an oriented graph and let $k \in \mathbb{N}$.

1. If $E = \emptyset$, then $\bar{\varphi}^*_G(k) = 1$.
2. If $e \in E$ is a bridge, then $\bar{\varphi}^*_G(k) = 0$.
3. If $e \in E$ is a loop, then $\bar{\varphi}^*_G(k) = (k+1) \cdot \bar{\varphi}^*_{G\backslash e}(k)$.
4. If $e \in E$ is neither a loop nor a bridge, then $\bar{\varphi}^*_G(k) = \bar{\varphi}^*_{G\backslash e}(k) + \bar{\varphi}^*_{G/e}(k)$.

To show this theorem we examine in how far a $\mathbb{Z}_k$-flow on G induces $\mathbb{Z}_k$-flows on G/e and $G\backslash e$, respectively, and in how far a totally cyclic orientation of G induces totally cyclic orientations of G/e and $G\backslash e$, respectively. We first turn our attention to the $\mathbb{Z}_k$-flows.

3.8.2. Lemma.

Let $G = (V, E)$ be an oriented graph and $e \in E$ neither a loop nor a bridge. If f is a $\mathbb{Z}_k$-flow on G, then $f|_{E\backslash e}$ is

1. a $\mathbb{Z}_k$-flow on G/e and
2. a $\mathbb{Z}_k$-flow on $G\backslash e$ if and only if $f(e) = 0$.

Proof. Let $e = uv$. At any vertex $w \notin \{u, v\}$ the flow $(Af)_w$ does not change when passing from G to G/e or $G\backslash e$. In G/e the vertices u and v have been identified to form a vertex u' and $(A_{G/e}f|_{E\backslash e})_{u'} = (Af)_u + (Af)_v = 0$. In $G\backslash e$ we have $(A_{G\backslash e}f|_{E\backslash e})_u = (Af)_u - f(e)$ which is zero if and only if $f(e) = 0$, and similarly for v. □

So a $\mathbb{Z}_k$-flow on G induces a $\mathbb{Z}_k$-flow on G/e and if $f(e) = 0$ it also induces a $\mathbb{Z}_k$-flow on $G\backslash e$. Moreover it turns out that any $\mathbb{Z}_k$-flow on G/e is induced by a unique $\mathbb{Z}_k$-flow on G and the same holds for $G\backslash e$.

3.8.3. Lemma.

Let $G = (V, E)$ be an oriented graph and $e \in E$ neither a loop nor a bridge.

1. Given a $\mathbb{Z}_k$-flow f' on G/e there is a unique $\mathbb{Z}_k$-flow f on G such that $f|_{E\backslash e} = f'$.
2. Given a $\mathbb{Z}_k$-flow f' on $G\backslash e$ there is a unique $\mathbb{Z}_k$-flow f on G such that $f|_{E\backslash e} = f'$. Moreover this flow has the property $f(e) = 0$.

Proof. In both cases, we necessarily have $f(e') = f'(e')$ for all $e' \neq e$ and we have to check that there is unique choice for $f(e)$ that makes f a $\mathbb{Z}_k$-flow. Let $e = uv$ oriented from u to

v. Let A^* denote the incidence matrix of G with the column corresponding to e removed. In both cases $(Af)_u = (A^*f)_u - f(e)$ and $(Af)_v = (A^*f)_v + f(e)$. So f is a $\mathbb{Z}_k$-flow if and only if $f(e) = (A^*f)_u$ and $f(e) = -(A^*f)_v$. In the first case these two values coincide because $(A^*f)_u + (A^*f)_v = (A_{G/e}f')_{u'} = 0$ where u' is the vertex obtained by identifying u and v. In the second case, both of these values are zero, because $(A^*f)_u = (A_{G\setminus e}f')_u = 0$ and $(A^*f)_v = (A_{G\setminus e}f')_v = 0$. □

In particular these two lemmas have a consequence that will be of importance to us later on.

3.8.4. Let $G = (V, E)$ be a graph and $S \subset E$ an edge set. The set of $\mathbb{Z}_k$-flows f on G with $\mathsf{zero}(f) \supset S$ is in bijection with the set of $\mathbb{Z}_k$-flows on $G \setminus S$.

A similar result can be shown for $\mathbb{Z}_k$-tensions.

3.8.5. Let $G = (V, E)$ be a graph and $S \subset E$ an edge set. The set of $\mathbb{Z}_k$-tensions t on G with $\mathsf{zero}(t) \supset S$ is in bijection with the set of $\mathbb{Z}_k$-tensions on G/S.

Note that in the case of flows we have to delete the set S while in the case of tensions we have to contract the set S. The statements become false if contract S in the flow case or delete S in the tension case.

Now we turn to totally cyclic orientations. Here the situation is a bit more complicated than with $\mathbb{Z}_k$-flows. We start with a useful characterization of totally cyclic orientations.

3.8.6. Lemma.

Let G be an oriented graph. σ is a totally cyclic orientation of G if and only if for any vertices $u, v \in V$ in the same component of the underlying undirected graph, there exists a directed path in ${}_\sigma G$ from u to v.

Proof. Suppose σ is totally cyclic. As both u and v lie in the same component of the undirected graph, there is an undirected path P from u to v. As ${}_\sigma G$ is totally cyclic, every edge of P lies on a directed cycle. In a directed cycle, there is a directed path from any vertex to any other vertex. So for any edge u_iv_i in P there is a directed path in G from u_i to v_i. Concatenating all these paths, we obtain a directed walk in G from u to v, which in particular contains a directed path from u to v as a subgraph.

Conversely, suppose we can always find a dipath from any vertex to any other. Let e be an edge oriented from u to v. Then the assumption guarantees the existence of a path P from v to u. Concatenating P and e yields a directed cycle. □

In the following, given an orientation σ and an edge e, we use the notation $\sigma \ominus e$ to denote the orientation obtained from σ by reversing the edge e.

3.8.7. Lemma.

Let $G = (V, E)$ be an oriented graph and $e \in E$ neither a loop nor a bridge. Let σ be a totally cyclic orientation of G. Then

1. $\sigma|_{E\setminus e}$ is a totally cyclic orientation of G/e, and
2. $\sigma|_{E\setminus e}$ is a totally cyclic orientation of $G\setminus e$ if and only if both σ and $\sigma \ominus e$ are totally cyclic orientations of G.

Proof. *(1)* As ${}_\sigma G$ is totally cyclic, there is a collection $\mathcal{C}$ of directed cycles in ${}_\sigma G$ that cover all edges. Then $\{C/e | C \in \mathcal{C}\}$ is a collection of directed cycles in ${}_{\sigma|_{E\setminus e}}G/e$ that covers all edges in G/e and hence ${}_{\sigma|_{E\setminus e}}G/e$ is totally cyclic.

(2) Let $e = uv$. Suppose ${}_{\sigma|_{E\setminus e}}G\setminus e$ is totally cyclic. Then by Lemma 3.8.6 there exist directed paths from u to v and from v to u in ${}_{\sigma|_{E\setminus e}}G\setminus e$. These show that no matter which way e is oriented in σ, we can always find a directed cycle in ${}_\sigma G$ on which e lies and so both ${}_\sigma G$ and ${}_{\sigma\ominus e}G$ are totally cyclic.

Conversely, suppose both ${}_\sigma G$ and ${}_{\sigma\ominus e}G$ are totally cyclic, where $\sigma(e) = +$. e lies on a directed cycle in ${}_{\sigma\ominus e}G$ so by Lemma 3.8.6 there is a directed path P from u to v in ${}_\sigma G\setminus e$. Let u', v' be any two vertices in $G\setminus e$. As ${}_\sigma G$ is totally cyclic there is a directed path P' in ${}_\sigma G$ from u' to v'. We replace every occurrence of e in P' with P and obtain a directed walk (and hence a directed path) in ${}_{\sigma|_{E\setminus e}}G\setminus e$ from u' to v'. By Lemma 3.8.6 it follows that ${}_{\sigma|_{E\setminus e}}G\setminus e$ is totally cyclic. □

3.8.8. Lemma.

Let $G = (V, E)$ be an oriented graph and $e \in E$ neither a loop nor a bridge.

1. Let $\sigma : E \setminus e \to \{+, -\}$ be a totally cyclic orientation of G/e. Then at least one of $\sigma \cup \{(e, +)\}$ and $\sigma \cup \{(e, -)\}$ is a totally cyclic orientation of G.
2. Let $\sigma : E \setminus e \to \{+, -\}$ be a totally cyclic orientation of $G\setminus e$. Then both $\sigma \cup \{(e, +)\}$ and $\sigma \cup \{(e, -)\}$ are totally cyclic orientations of G.

Proof. *(1)* Let ${}_\sigma G/e$ be totally cyclic and $e = uv$. Let $\mathcal{C}$ be a collection of directed cycles in ${}_\sigma G/e$ that covers all edges of G/e. Now we distinguish two cases: Is one of these cycles "broken" in G or not? More precisely does there exist a cycle $C \in \mathcal{C}$ that contains consecutive edges e_1 and e_2 such that e_1 enters u and e_2 leaves v (or vice versa)? If not, then $\mathcal{C}$ shows that ${}_\sigma G\setminus e$ is also totally cyclic and we can continue as in part *(2)* below.

So we suppose that C is such a broken cycle. In this case C gives a directed path from v to u in G. We now orient e from u to v. Then any directed path P in G/e from a vertex u' to a vertex v' can be turned into a directed path in G from u' to v' by substituting the edge e or the path given by C wherever P is broken. Using Lemma 3.8.6 the claim follows.

(2) Already in $G\backslash e$ there is for any two vertices u, v in the same component a directed path from u to v. This remains true after the edge e is inserted, no matter how e is oriented (note that e is not a bridge). So by Lemma 3.8.6 both $\sigma \cup \{(e,+)\}$ and $\sigma \cup \{(e,-)\}$ are totally cyclic orientations of G. □

Now we have all the pieces together to show that $\bar{\varphi}^*_G(k)$ is a Tutte-Grothendieck invariant.

Proof of Theorem 3.8.1. *1.* If $E = \emptyset$, then $\mathsf{F}_G(k) = \{(\emptyset, \emptyset)\}$.

2. If $e \in E$ is a bridge, then any flow f on G has $f(e) = 0$. Thus e is also a bridge in $G/\mathsf{supp}(f)$ and which means that there is no totally cyclic orientation on $G/\mathsf{supp}(f)$. So $\mathsf{F}_G(k) = \emptyset$.

3. If $e \in E$ is a loop, then $(f, \sigma) \mapsto (f|_{E\backslash e}, \sigma \cap E \setminus e)$ is a surjective map from $\mathsf{F}_G(k)$ onto $\mathsf{F}_{G\backslash e}$ and every fiber of this map has cardinality $k+1$. The reason is that given $(f|_{E\backslash e}, \sigma \cap E \setminus e)$ we can define $f(e) \in \mathbb{Z}_k$ arbitrarily and f will become a $\mathbb{Z}_k$-flow on G. The case $f(e) = 0$ is counted twice as either orientation of e will turn σ into a totally cyclic orientation of $G/\mathsf{supp}(f)$.

4. Let $e \in E$ be neither a bridge nor a loop. Consider the map $\pi_{G/e} : \mathsf{F}_G(k) \to \mathsf{F}_{G/e}(k)$ given by $(f, \sigma) \mapsto (f|_{E\backslash e}, \sigma|_{E\backslash e})$. Lemmas 3.8.2 and 3.8.7 tell us that $\pi_{G/e}$ is well-defined and Lemmas 3.8.3 and 3.8.8 tell us that every $(f', \sigma') \in \mathsf{F}_{G/e}(k)$ has either one or two preimages under $\pi_{G/e}$. (f', σ') has two preimages if and only if the unique $\mathbb{Z}_k$-flow f with $f|_{E\backslash e} = f'$ has $f(e) = 0$ and both $\sigma' \cup \{(e,+)\}$ and $\sigma' \cup \{(e,-)\}$ are totally cyclic orientations of $G/\mathsf{supp}(f)$.

Loosely speaking, this means that the cardinalities of $\mathsf{F}_G(k)$ and $\mathsf{F}_{G/e}(k)$ are the same, except that we have to count those $(f', \sigma') \in \mathsf{F}_{G/e}(k)$ that have two preimages twice.

So let $\mathsf{F}'_G(k)$ denote the set of all $(f, \sigma) \in \mathsf{F}_G(k)$ such that $f(e) = 0$ and both σ and $\sigma \ominus e$ are totally cyclic orientations on $G/\mathsf{supp} f$. Consider the map $\pi_{G\backslash e} : \mathsf{F}'_G(k) \to \mathsf{F}_{G\backslash e}(k)$ given by $(f, \sigma) \mapsto (f|_{E\backslash e}, \sigma \cap E \setminus e)$. Lemmas 3.8.2 and 3.8.7 tell us that $\pi_{G\backslash e}$ is well-defined and Lemmas 3.8.3 and 3.8.8 tell us that every $(f', \sigma') \in F_{G\backslash e}(k)$ has precisely two preimages under $\pi_{G\backslash e}$. But this means that $\bar{\varphi}^*_G(k) = \bar{\varphi}^*_{G/e}(k) + \bar{\varphi}^*_{G\backslash e}(k)$ as desired. □

3.9. *A Geometric View on Modular Tension Reciprocity*

We already gave a proof of the Modular Tension Reciprocity Theorem 3.3.6 by showing its equivalence to Stanley's Chromatic Reciprocity Theorem in Section 3.3. In this section we give a brief outline how the reciprocity theorem for modular tensions can be obtained from geometric considerations in the spirit of Section 3.6 and how the modular tension polynomial can be obtained as the Ehrhart polynomial of an inside-out polytope in the spirit of Section 3.7. As we found out after completing our proofs, Chen [Che07] employs a similar geometric approach to show a reciprocity theorem for the modular tension polynomial.

First, we sketch how the argument from Section 3.6 can be transferred to the case of modular tensions. Let $G = (V, E)$ be a graph. Again we identify $\mathbb{Z}_k$ with the set of integers $0, \ldots, k-1$, which allows us to identify nowhere zero maps $t : E \to \mathbb{Z}_k$ with lattice points in $(0, k)^E$. Let

T be a spanning multi-tree of G and C the corresponding cycle matrix. The $\mathbb{Z}_k$-tensions then correspond to those $t \in \mathbb{Z}^E \cap (0,k)^E$ with $Ct \equiv 0 \mod k$, that is, those t for which there exists a $d \in \mathbb{Z}^{|E|-|V|+c}$ such that $Ct = kd$. We then define the open and closed polytopes

$$\begin{aligned} P_d^\circ &= \{t \in \mathbb{Z}^E | 0 < t < 1, Ct = d\} \\ P_d &= \{t \in \mathbb{Z}^E | 0 \leq t \leq 1, Ct = d\} \end{aligned}$$

and call those $d \in \mathbb{Z}^{|E|-|V|+c}$ for which $P_d^\circ \neq \emptyset$ **feasible**. The set of all feasible d we denote by $\mathcal{D}_G$. Then the modular tension polynomial is the sum of Ehrhart functions

3.9.1. $\bar{\theta}_G(k) = \sum_{d \in \mathcal{D}_G} L_{P_d^\circ}(k) = \# \bigcup_{d \in \mathcal{D}_G} \mathbb{Z}^E \cap kP_d^\circ.$

If d is feasible $P_d^\circ = \operatorname{relint} P_d$ and the dimension of the P_d° is the dimension of the tension space $\dim P_d^\circ = |V| - c$. Moreover the P_d are lattice polytopes. Applying Ehrhart-Macdonald Reciprocity we obtain

3.9.2. $(-1)^{|V|-c}\bar{\theta}_G(-k) = \sum_{d \in \mathcal{D}_G} L_{P_d}(k) = \# \bigcup_{d \in \mathcal{D}_G} \mathbb{Z}^E \cap kP_d.$

Using 3.2.7 and a similar argumentation as in Section 3.6 we can interpret the elements of $\bigcup_{d \in \mathcal{D}_G} \mathbb{Z}^E \cap kP_d$ as pairs (t, σ) of a $\mathbb{Z}_k$-tension t on G and an acyclic orientation σ of $G \setminus \operatorname{supp}(t)$. Note that this time we have to delete the edges in $\operatorname{supp}(t)$ instead of contracting them. All in all this gives a geometric proof of the Modular Tension Reciprocity Theorem 3.3.6.

Second, we explain how the construction in Section 3.7 can be applied to obtain an inside-out polytope whose Ehrhart function is the modular tension polynomial. Again, let T denote a spanning multi-tree of our graph $G = (V, E)$ and let N^t be the transpose of the corresponding network matrix. The columns of N^t are indexed by the edges of T and the rows are indexed by the edges in E such that the edges in T come first. The row $N^t_{e\cdot}$ of N^t indexed by $e = uv \in E$ is the sign vector of the path in T from u to v. So given a map $z : T \to \mathbb{R}$, the value $\langle N^t_{e\cdot}, z \rangle$ is the weight e would have to have in an $\mathbb{R}$-tension t with $t|_T = z$. So the image of the open cube $(0,k)^T$ under the matrix N^t is precisely the set of $\mathbb{R}$-tensions t with $0 < t(e) < k$ for all $e \in T$. We can obtain the set of $\mathbb{R}$-tensions t with $0 < t(e) < k$ for all $e \in T$ and $t(e) \not\equiv 0 \mod k$ for all $e \in E$ in the image by excluding those points from $(0,k)^T$ that lie on one of the hyperplanes $\langle N^t_{e\cdot}, z \rangle = kd_e$ for $e \in E$ and $d_e \in \mathbb{Z}$. Just as in the case of flows it turns out that the lattice points lying in this image are in bijection with the points in $\bigcup_{d \in \mathcal{D}_G} \mathbb{Z}^E \cap kP_d^\circ$ and consequently with the nowhere zero $\mathbb{Z}_k$-tensions on G. So if we define $\mathcal{H}$ to be the hyperplane arrangements consisting of those hyperplanes $\{z | \langle N^t_{e\cdot}, z \rangle = d_e\}$ for $e \in E$ and $d_e \in \mathbb{Z}$ that meet the open cube $(0,1)^T$, then

3.9.3. $\bar{\theta}_G(k) = \mathsf{L}_{(0,1)^T \setminus \bigcup \mathcal{H}}(k)$.

Again, the closed cells of the inside-out polytope $((0,1)^T, \mathcal{H})$ are lattice polytopes, because N is totally unimodular, and they are lattice transforms of the polytopes P_d for $d \in \mathcal{D}_G$.

3.10. $\mathcal{B}_G$ and $\mathcal{D}_G$

The sets $\mathcal{B}_G$ and $\mathcal{D}_G$ have played a crucial role in our geometric proofs of the Modular Flow and Tension Reciprocity Theorems. Yet, we have not said much about them. In this section we will investigate them more closely, starting with $\mathcal{B}_G$. The benefit will be combinatorial interpretations of $\bar{\varphi}_G(0)$ and $\bar{\theta}_G(0)$.

$|\mathcal{B}_G|$ is the number of non-empty open polytopes P_b° that arise by intersecting an integral translate of the flow space with the unit cube $[-1,1]^E$. Their closures P_b are lattice polytopes. In 3.6.6 we have seen that $(-1)^{|E|-|V|+c} \cdot \bar{\varphi}(-k) = \sum_{b \in \mathcal{B}_G} \mathsf{L}_{P_b}(k)$ and this holds as an identity of polynomials for all k. The Ehrhart function of a lattice polytope evaluated at 0 is 1. Thus

3.10.1. $(-1)^{|E|-|V|+c} \bar{\varphi}_G(0) = |\mathcal{B}_G|$.

So interpreting $|\mathcal{B}_G|$ amounts to finding a combinatorial interpretation of $\bar{\varphi}_G(0)$. Such interpretations are known [Gio07],[BO92],[EMM08], but what does the geometric approach lead to?

$\bigcup_{b \in \mathcal{B}_G} P_b$ forms a partition of a subset of the vertices of the unit cube $[0,1]^E$. Which $v \in \{0,1\}^E$ lie in P_b for some $b \in \mathcal{B}_G$? Our work in Section 3.6 implies that these v are characterized by the property that there exists a nowhere zero $\mathbb{R}$-flow f such that $v(e) = 0$ iff $\mathsf{sgn}(f(e)) = +$ and $v(e) = 1$ iff $\mathsf{sgn}(f(e)) = -$. So by 3.2.3 the vertices $\{0,1\}^E \cap \bigcup_{b \in \mathcal{B}_G} P_b$ are in bijection with the totally cyclic orientations of G. The vertex v and the corresponding totally cyclic orientation σ are related by $\sigma = 1 - 2v$, where 1 denotes the all-one vector. The P_b partition these vertices into classes. For two totally cyclic orientations σ_1 and σ_2 we write $\sigma_1 \sim \sigma_2$ if there exists a $b \in \mathcal{B}_G$ such that the vertices $v_1, v_2 \in \{0,1\}^E$ corresponding to σ_1 and σ_2, respectively, both lie in P_b. In this case we call the two **equivalent**.

Two vertices v_1 and v_2 lie in the same P_b if and only if $v_2 - v_1 \in \ker A$. In other words:

3.10.2. $v_1 \sim v_2$ if and only if $v_2 - v_1$ is a $0, \pm 1$-flow.

On the other hand, given a $0, \pm 1$-flow f and a vertex v of P_b, then $v + f$ is a vertex of P_B if and only if for all $e \in E$

$$v(e) = 0 \Rightarrow f(e) \geq 0 \text{ and } v(e) = 1 \Rightarrow f(e) \leq 0. \tag{3.6}$$

This leads to the following criterion.

Let σ be an orientation of $G = (V, E)$. Building on the notation given in Section 3.8 we define $\sigma \ominus S$ to be the orientation obtained from σ by reversing all edges in $S \subset E$. That is $(\sigma \ominus S)|_{E\setminus S} = \sigma|_{E\setminus S}$ and $(\sigma \ominus S)|_S = -\sigma|_S$. This operation we call the **reversal** of S. By Lemma 3.8.6 we already know:

3.10.3. Let C be a directed cycle in $_\sigma G$. Then σ is a totally cyclic orientation of G if and only if $\sigma \ominus C$ is a totally cyclic orientation of G.

But this can also be seen from our above considerations. If σ corresponds to a vertex v of P_b, then a directed cycle C corresponds to the $0, \pm 1$-flow f with $\mathsf{supp}(f) = C$, $f(e) = 1$ if $v(e) = 0$ and $f(e) = -1$ if $v(e) = 1$. The orientation $\sigma \ominus C$ then corresponds to the vertex $v + f$. If C is a directed cycle in $_\sigma G$, then σ and $\sigma \ominus C$ both correspond to vertices of the same P_b. We say that an orientation σ_2 can be obtained from an orientation σ_1 by **reversals of directed cycles**, if there exists a sequence of cycles $C_1, \ldots, C_l$ in G such that

1. C_i is a directed cycle in $\sigma_1 \ominus C_1 \ominus \ldots \ominus C_{i-1}$ for all $1 \le i \le l$, and
2. $\sigma_2 = \sigma_1 \ominus C_1 \ominus \ldots \ominus C_l$.

Then we can characterize equivalent totally cyclic orientations as follows.

3.10.4. Totally cyclic orientations σ_1 and σ_2 are equivalent if and only if σ_2 can be obtained from σ_1 by the reversal of directed cycles.

Proof. The above arguments show that if σ_2 can be obtained from σ_1 by the reversal of directed cycles, then the two are equivalent. To show the converse we argue by induction on the number of edges that are oriented differently by σ_1 and σ_2. Let v_1 and v_2 be the corresponding vertices and let S be the set of edges where σ_1 and σ_2 differ. Clearly $\mathsf{supp}(\sigma_1 - \sigma_2) = S$. But $\frac{1}{2}(\sigma_1 - \sigma_2) = v_2 - v_1 =: f$ is a $0, \pm 1$-flow and so $\mathsf{sgn}(f)|_S$ is a totally cyclic orientation of $G[S]$. On the other hand $\mathsf{sgn}(f)|_S = \sigma_1|_S$ and so $_{\sigma_1}G[\mathsf{supp}(f)]$ is totally cyclic and thus contains a directed cycle C. Therefore $\sigma_1 \ominus C$ is a totally cyclic orientation of G that is equivalent to σ_1 and $\sigma_1 \ominus C$ and σ_2 differ on strictly fewer edges than do σ_1 and σ_2. The claim follows by induction. □

What is the significance of the fact that we are only allowed to reverse directed cycles and not arbitrary cycles?

Reversing a *directed* cycle does not change the sequence of in-degrees.

The **in-degree** at a vertex v is the number of in-edges at v in $_\sigma G$, which we denote by $I_\sigma(v)$. Similarly the **out-degree** at a vertex v is the number of out-edges at v in $_\sigma G$, which we denote by $O_\sigma(v)$. The **degree** at v is $D(v) = I_\sigma(v) + O_\sigma(v)$, which is independent of σ. Thus for a fixed graph G the **in-degree sequence** I_σ of a given orientation σ determines its **out-degree sequence** O_σ and vice versa.

Now we make the following simple observation: If A is the incidence matrix of G and σ an orientation, then

3.10.5. $A\sigma = I_\sigma - O_\sigma$

and again, because the degree sequence is fixed, $A\sigma$ determines both I_σ and O_σ. Now consider two totally cyclic orientations σ_1 and σ_2 and their corresponding vertices v_1 and v_2. We have $A\sigma_1 = A\sigma_2$ if and only if

$$0 = A(\sigma_1 - \sigma_2) = 2A(v_2 - v_1),$$

that is if and only if v_1 and v_2 are vertices of the same polytope P_b. So this proves that

3.10.6. $\mathcal{B}_G$ is in bijection with the set of in-degree sequences of totally cyclic orientations of G,

which implies the following recent theorem of Gioan [Gio07] as a corollary. Here $\bar{\varphi}^*_G$ denotes the reciprocal of the modular flow polynomial $\bar{\varphi}_G$.

3.10.7. Theorem. (Gioan [Gio07, Theorem 3.1])

$(-1)^{|E|-|V|+c}\bar{\varphi}_G(0) = \bar{\varphi}^*_G(0)$ is the number of in-degree sequences of totally cyclic orientations of G.

Gioan [Gio07] also studies "cycle reversal systems". So this section does not present new results, but rather shows how these concepts arise naturally from the geometric setup.

What we did above for the modular flows can also be achieved for the modular tensions. To interpret $\mathcal{D}_G$ we proceed exactly as in our interpretation of $\mathcal{B}_G$, so we do not give all the details but only state the key facts. First of all the vertices v of $[0,1]^E$ contained in $\bigcup_{d \in \mathcal{D}_G} P_d$ correspond precisely to the acyclic orientations of σ of G via $\sigma = 1 - 2v$. Again, we say that two acyclic orientations σ_1, σ_2 are equivalent if the corresponding vertices v_1, v_2 lie in the same P_d for $d \in \mathcal{D}_G$, that is if $v_2 - v_1$ is a $0, \pm 1$-tension. The role of directed cycles is played, this time, by the signed or directed cuts. Note that the $0, \pm 1$-tension $v_2 - v_1$ itself is not necessarily a directed cut but it can be written as the sum of directed cuts with disjoint support. Thus we can find cuts $C_1, \ldots, C_l$ such that $\sigma_1 \ominus C_1 \ominus \ldots \ominus C_l = \sigma_2$ and C_i is a directed cut in $\sigma_1 \ominus C_1 \ominus \ldots \ominus C_{i-1}$ for all i. This leads to the characterization:

3.10.8. Acyclic orientations σ_1 and σ_2 are equivalent if and only if σ_2 can be obtained from σ_1 by the reversal of directed cuts.

Again Gioan [Gio07] also studies "cocycle reversal systems", but he does not draw a connection to degrees. So let us briefly consider the question: What concept of "degree" does the reversal of directed cuts leave invariant? The in-edges at a vertex v can be viewed as the edges going from the shore $E \setminus v$ to the shore v of the single vertex cut $E \setminus v \cup v$. If we call the edges going from shore X to shore Y the in-edges of a cut $X \cup Y$ then reversing the edges along a directed cycle does not change the number of in-edges along any cut. So to come to a dual concept of in-degree, we can switch the roles of cycles and cuts.

Let C be a cycle in G along with an orientation σ_C such that ${}_{\sigma_C}G[C]$ is a directed cycle. Let σ be any orientation of G. The **in-codegree** of σ on the cycle C is the number of edges $e \in E$ such that $\sigma(e) = \sigma_C(e)$, while the **out-codegree** of σ on C is the number of edges with $\sigma(e) = -\sigma_C(e)$. Again in-codegree and out-codegree determine one another and the reversal of a directed cut does not change the in-codegree on any cycle. The **in-codegree sequence** of σ is simply the function I^*_σ that assigns to any pair C, σ_C the in-codegree of σ on C with respect to σ_C and we define the **out-codegree sequence** O^*_σ similarly. Note that $I_\sigma - O_\sigma$ is just the image of σ under a suitably chosen cycle matrix of G. This shows that

3.10.9. $\mathcal{D}_G$ is in bijection with the set of in-codegree sequences of acyclic orientations of G.

Consequently:

3.10.10. Theorem.

$(-1)^{|V|-c}\bar{\theta}_G(0) = \bar{\theta}^*_G(0)$ is the number of in-codegree sequences of acyclic orientations of G.

3.11. The Tutte Polynomial as a Counting Function

In this section we show how the reciprocity results for modular tensions and flows connect a convolution formula for the Tutte polynomial T_G to a beautiful interpretation of T_G as a counting function. This interpretation appears in the guise of a formula in Reiner's article [Rei99, Corollary 9]. In this section we give a proof of the equivalence of the counting interpretation and the convolution formula that follows naturally from the reciprocity theorems for modular flows and tensions. Putting these results together with Theorems 3.10.7 and 3.10.10 we obtain a combinatorial interpretation of the value of the Tutte polynomial at every lattice point in the plane.

Given a graph $G = (V, E)$ the Tutte polynomial T_G is the polynomial in two variables that is defined by the following deletion/contraction formula.

1. If $E = \emptyset$, then $T_G(x, y) = 1$.
2. If $e \in E$ is a bridge, then $T_G(x, y) = x \cdot T_{G/e}(x, y)$.
3. If $e \in E$ is a loop, then $T_G(x, y) = y \cdot T_{G \setminus e}(x, y)$.
4. If $e \in E$ is neither a loop nor a bridge, then $T_G(x, y) = T_{G/e}(x, y) + T_{G \setminus e}(x, y)$.

A great deal more can be said about the Tutte polynomial; we refer the interested reader to [GR04],[Aig07],[EMM08] and [BO92]. Most importantly it is known that all Tutte-Grothendieck invariants can be expressed in terms of the Tutte polynomial. This also applies to the modular flow and tension polynomials.

3.11.1. Theorem. (Tutte [Tut54])

Let $G = (V, E)$ be a graph with c components. Then

$$\begin{aligned}(-1)^{|V|-c} \cdot \bar{\theta}_G(x) &= T_G(1-x, 0), \\ (-1)^{|E|-|V|+c} \cdot \bar{\varphi}_G(y) &= T_G(0, 1-y).\end{aligned}$$

In light of the Modular Flow and Tension Reciprocity Theorems 3.3.5 and 3.3.6 these identities are striking! They simply state that $T_G(1+x, 0)$ is the reciprocal of the modular tension polynomial and that $T_G(0, 1+y)$ is the reciprocal of the modular flow polynomial. Letting $\bar{\theta}^*_G(x) = (-1)^{|V|-c}\bar{\theta}_G(x)$ and $\bar{\varphi}^*_G(y) = (-1)^{|E|-|V|+c}\bar{\varphi}_G(y)$, respectively, denote the reciprocals of the modular tension and flow polynomials, we have

3.11.2.

$$\bar{\theta}^*_G(x) = T_G(1+x, 0) \quad \text{and} \quad \bar{\varphi}^*_G(y) = T_G(0, 1+y).$$

This can be combined with the following convolution formula by Kook, Reiner and Stanton [KRS99]

3.11.3. Convolution Formlua.

$$T_G(x, y) = \sum_{S \subset E} T_{G/S}(x, 0) T_{G[S]}(0, y)$$

to yield

3.11.4.

$$T_G(1+x, 1+y) = \sum_{S \subset E} \bar{\theta}^*_{G/S}(x) \cdot \bar{\varphi}^*_{G[S]}(y).$$

By the Modular Flow and Tension Reciprocity Theorems 3.3.5 and 3.3.6 we can now interpret this sum.

3.11.5. Theorem.

Let $G = (V, E)$ be a graph with c components. Then for all $0 < x, y \in \mathbb{N}$

$$\begin{aligned} T_G(1+x, 1+y) = \#\{ (S, t, f) \; : \; & S \subset E, \\ & t \text{ a } \mathbb{Z}_x\text{-tension on } G/S, \\ & f \text{ a } \mathbb{Z}_y\text{-flow on } G[S] \; \}. \end{aligned}$$

Theorem 3.11.5 is a rephrasing of Reiner [Rei99, Corollary 9]. Note that in [Rei99] it is not clear whether this result is stated for all positive x, y or only for prime powers. We show Theorem 3.11.5 for all positive x, y. In this section we show the equivalence of Theorem 3.11.5 to the Convolution Formula using the Modular Reciprocity Theorems. In the next section we give an inductive proof of Theorem 3.11.5 from first principles.

To our knowledge Theorem 3.11.5 is the only interpretation of $T_G(1+x, 1+y)$ as a counting function that holds at all lattice points $(x, y) \in \mathbb{Z}^2_{\geq 1}$ in the positive quadrant. Note that $T_G(1+x, 1+y)$ cannot in general be captured as a counting function in the other three quadrants, as for many graphs $T_G(1+x, 1+y)$ has both positive and negative values in each of the other three quadrants. [Rei99, Corollary 2] gives a "combinatorial interpretations" of T_G in the negative quadrant, that does not interpret T_G simply as a counting function but instead expresses T_G as a sum in which both positive and negative summands appear. Using the reciprocity theorems it is also possible to obtain similar interpretations for the remaining two quadrants as we will do at the end of this section.

In light of 3.8.4 and 3.8.5, several variations of Theorem 3.11.5 suggest themselves. One is to view $T_G(1+x, 1+y)$ as counting pairs (t, f) of a $\mathbb{Z}_x$-tension t and a $\mathbb{Z}_y$-flow f on G with disjoint support, where each pair is counted with multiplicity $2^{|E \setminus \mathsf{supp}(t) \cup \mathsf{supp}(f)|}$. This can be written using the formula

$$T_G(1+x, 1+y) = \sum_{\substack{t \text{ a } \mathbb{Z}_x\text{-tension on } G \\ f \text{ a } \mathbb{Z}_y\text{-flow on } G \\ \mathsf{supp}(t) \cap \mathsf{supp}(f) = \emptyset}} 2^{|E \setminus \mathsf{supp}(t) \cup \mathsf{supp}(f)|}$$

which is essentially the way this result appears in [Rei99]. Another way of putting this is to write

$$T_G(1+x, 1+y) = \sum_{\substack{S,T \subset E \\ S \cap T = \emptyset}} 2^{|E \setminus S \setminus T|} \cdot \bar{\theta}_{G/(E \setminus T)}(x) \cdot \bar{\varphi}_{G[S]}(y)$$

which complements 3.11.4. It may however be worthwhile to interpret the multiplicity $2^{|E \setminus \mathsf{supp}(t) \cup \mathsf{supp}(f)|}$ as counting the number of (arbitrary) orientations σ of $G \setminus \mathsf{supp}(t) / \mathsf{supp}(f)$ and write

$$\begin{aligned} T_G(1+x, 1+y) \quad = \quad \#\{\ (t, f, \sigma) \quad : \quad & t \text{ a } \mathbb{Z}_x\text{-tension on } G, \\ & f \text{ a } \mathbb{Z}_y\text{-flow on } G, \\ & \mathsf{supp}(t) \cap \mathsf{supp}(f) = \emptyset \\ & \sigma \text{ an orientation of } G \setminus \mathsf{supp}(t)/\mathsf{supp}(f)\}. \end{aligned} \tag{3.7}$$

The reason for this is the following lemma, which is crucial for our proof of the equivalence of Theorems 3.11.3 and 3.11.5.

3.11.6. Lemma.

Let $G = (V, E)$ be an oriented graph. Then there is a unique edge set $S \subset E$ such that $G[S]$ is totally cyclic and G/S is acyclic.

Proof. *Existence of S.* Let S be the set of edges in E that lie on a directed cycle in G. Then $G[S]$ is totally cyclic. Moreover, in every component of K of $G[S]$ there is a directed path between any two vertices in K. Thus, an edge e does not lie on a directed cycle in G/S if and only if e does not lie on a directed cycle in G. So G/S is acyclic as desired.

Uniqueness of S. Now suppose $S' \subset E$ is any edge set such that $G[S']$ is totally cyclic and G/S' is acyclic. Since $G[S']$ is totally cyclic, $S' \subset S$. Now suppose there exists an edge $e \in S \setminus S'$. Then e is contained in a directed cycle in G. But this remains true in G/S as contracting edges does not destroy directed cycles. So G/S is not acyclic which is a contradiction. Thus $S = S'$. □

Now we can show the equivalence of the Convolution Formula and our interpretation of $T_G(1+x, 1+y)$ for positive x and y.

Equivalence of 3.11.3 and 3.11.5. We have already seen that the Convolution Formula 3.11.3 is equivalent to 3.11.4, which simply states that $T_G(1+x, 1+y)$ counts the number of tuples $(S, t, \sigma_t, f, \sigma_f)$ of a set $S \subset E$, a $\mathbb{Z}_x$-tension t on G/S, an acyclic orientation σ_t of $G/S \setminus \mathsf{supp}(t)$, a $\mathbb{Z}_y$-flow f on $G[S]$ and a totally cyclic orientation σ_f of $G[S]/\mathsf{supp}(f)$.

t induces a $\mathbb{Z}_x$-tension t' on G with $S \subset \mathsf{zero}(t)$ and $t'|_{E\setminus S} = t$ and f induces a $\mathbb{Z}_y$-flow f' on G with $E \setminus S \subset \mathsf{zero}(f)$ and $f'|_S = f$, while $\sigma := \sigma_t \cup \sigma_f$ is an orientation of $G \setminus \mathsf{supp}(t)/\mathsf{supp}(f)$. This defines a map $h : (S, t, \sigma_t, f, \sigma_f) \mapsto (t', f', \sigma)$.

We now construct the inverse map $g : (t', f', \sigma) \mapsto (S, t, \sigma_t, f, \sigma_f)$. For this purpose we use the abbreviation $G' := G \setminus \mathsf{supp}(t)/\mathsf{supp}(f)$ and let E' denote the edge set of G'.

Given a triple (t', f', σ) consisting of a $\mathbb{Z}_x$-tension t' on G and a $\mathbb{Z}_y$-flow f' on G with disjoint support, as well as an orientation σ of G' we can define $(S, t, \sigma_t, f, \sigma_f)$ as follows. By Lemma 3.11.6, we know that there is a unique set $S' \subset E'$ such that $\sigma|_{S'} =: \sigma_f$ is a totally cyclic orientation of $G'[S']$ and $\sigma|_{E'\setminus S'} =: \sigma_t$ is an acyclic orientation of G'/S'. We define $S := \mathsf{supp}(f) \cup S'$, $t := t'|_{E\setminus S}$ and $f := f'|_S$. This completes the definition of our map g, which is the inverse of h. Therefore (3.7) is equivalent to the Convolution Formula.

To see that Theorem 3.11.5 is equivalent to the Convolution Formula, we note that given (t', f', σ), we can define $S := S' \cup \mathsf{supp}(f)$ where S' consists of those edges $e \in E'$ with $\sigma(e) = +$ and put $t := t'|_{E\setminus S}$ and $f := f'|_S$. This defines the required bijection $(t', f', \sigma) \mapsto (S, t, f)$. □

But we can go further from here. The Modular Flow and Tension Reciprocity Theorems together with Theorems 3.10.7 and 3.10.10 provide us with a way of interpreting the values of $\bar{\varphi}_G(k)$ and $\bar{\theta}_G(k)$ at every integer k. And the Convolution Formula tells us that we only need to know the values of $\bar{\varphi}_G$ and $\bar{\theta}_G$ if we want to find out what a particular value of the Tutte polynomial is.

> The Convolution Formula together with the Modular Reciprocity Theorems give a *unified framework* for interpreting *all values* of the Tutte polynomial.

This leads us to the following theorem, that gives an interpretation of $T_G(1+x, 1+y)$ as a counting function for every *non-negative* integers x and y and an interpretation of $T_G(1+x, 1+y)$ as the difference of two counting functions at all other integers x and y. As explained above this is best possible in the sense that the Tutte polynomial takes both negative and positive values in all quadrants of the plane except the positive one. To our knowledge this is the only interpretation of the Tutte polynomial of this generality. Certainly there are "nicer" interpretations for some values of the Tutte polynomial than the ones we give below, but the advantage of our approach is that we arrive at all of these interpretations using a single method.

Part 1. of the following theorem is Theorem 3.11.5 while part 5. is a rephrasing of [Rei99, Corollary 2].

3.11.7. Theorem.

Let $G = (V, E)$ be a graph with c components. Let $0 < x, y \in \mathbb{N}$. Then the following identities hold. In the positive quadrant we can interpret T_G as a counting function.

1. $T_G(1+x, 1+y) = \#\{ (S, t, f) :$
   $S \subset E$,
   t a $\mathbb{Z}_x$-tension on G/S,
   f a $\mathbb{Z}_y$-flow on $G[S]$ $\}$.

2. $T_G(1+x, 1) = \#\{ (S, t, \sigma, I) :$
   $S \subset E$,
   t a $\mathbb{Z}_x$-tension on G/S,
   σ an acyclic orientation of $G/S \setminus \mathsf{supp}(t)$,
   I an in-degree sequence of a totally cyclic orientation of $G[S]$ $\}$.

3. $T_G(1, 1+y) = \#\{ (S, I^*, f, \sigma) :$
   $S \subset E$,
   I^* an in-codegree sequence of an acyclic orientation of G/S,
   f a $\mathbb{Z}_y$-flow on $G[S]$,
   σ a totally cyclic orientation of $G[S]/\mathsf{supp}(f)$ $\}$.

4. $T_G(1, 1) = \#\{ (S, I^*, I) :$
   $S \subset E$,
   I^* an in-codegree sequence of an acyclic orientation of G/S,
   I an in-degree sequence of a totally cyclic orientation of $G[S]$ $\}$.

Everywhere else we can interpret the Tutte polynomial, up to sign, as a difference of two counting functions.

5. $(-1)^{|V|+c} \cdot T_G(1-x, 1-y) = + \#\{(S,t,f) \text{ such that } |S| \text{ is even}\} - \#\{(S,t,f) \text{ such that } |S| \text{ is odd}\}$

where $S \subset E$,
t a nowhere zero $\mathbb{Z}_x$-tension on G/S,
f a nowhere zero $\mathbb{Z}_y$-flow on $G[S]$.

6. $(-1)^{c} \cdot T_G(1-x, 1+y) = + \#\{(S,t,f,\sigma) \text{ such that } c(G[S]) \text{ is even}\} - \#\{(S,t,f,\sigma) \text{ such that } c(G[S]) \text{ is odd}\}$

where $S \subset E$,
t a nowhere zero $\mathbb{Z}_x$-tension on G/S,
f a $\mathbb{Z}_y$-flow on $G[S]$,
σ a totally cyclic orientation of $G[S]/\mathsf{supp}(f)$.

7. $(-1)^{|V|} \cdot T_G(1+x, 1-y) = + \#\{(S,t,\sigma,f) \text{ such that } |S|+c(G[S]) \text{ is even}\} - \#\{(S,t,\sigma,f) \text{ such that } |S|+c(G[S]) \text{ is odd}\}$

where $S \subset E$,
t a $\mathbb{Z}_x$-tension on G/S,
σ an acyclic orientation of $G/S \setminus \mathsf{supp}(t)$
f a nowhere zero $\mathbb{Z}_y$-flow on $G[S]$.

8. $(-1)^{c} \cdot T_G(1-x, 1) = + \#\{(S,t,I) \text{ such that } c(G[S]) \text{ is even}\} - \#\{(S,t,I) \text{ such that } c(G[S]) \text{ is odd}\}$

where $S \subset E$,
t a nowhere zero $\mathbb{Z}_x$-tension on G/S,
I an in-degree sequence of a totally cyclic orientation of $G[S]$.

9. $(-1)^{|V|} \cdot T_G(1, 1-y) = + \#\{(S,I^*,f) \text{ such that } |S|+c(G[S]) \text{ is even}\} - \#\{(S,I^*,f) \text{ such that } |S|+c(G[S]) \text{ is odd}\}$

where $S \subset E$,
I^* an in-codegree sequence of an acyclic orientation of G/S,
f a nowhere zero $\mathbb{Z}_y$-flow on $G[S]$.

Proof. *1.* This is Theorem 3.11.5.

2. $T_G(1+x, 1) = \sum_{S \subset E} \bar{\theta}^*_{G/S}(x) \cdot \bar{\varphi}^*_{G[S]}(0)$. By Theorem 3.3.6, $\bar{\theta}^*_{G/S}(x)$ is the number of pairs (t, σ) of a $\mathbb{Z}_x$-tension on G/S and an acyclic orientation σ of $G/S \setminus \mathsf{supp}(t)$. By Theorem 3.10.10, $\bar{\varphi}^*_{G[S]}(0)$ is the number of in-degree sequences of totally cyclic orientations of $G[S]$.

3. $T_G(1, 1+y) = \sum_{S \subset E} \bar{\theta}^*_{G/S}(0) \cdot \bar{\varphi}^*_{G[S]}(y)$. By Theorem 3.10.10, $\bar{\theta}^*_{G/S}(0)$ is the number of in-codegree sequences of acyclic orientations of G/S. By Theorem 3.3.5, $\bar{\varphi}^*_{G[S]}(y)$ is the number of pairs (f, σ) of a $\mathbb{Z}_y$-flow on $G[S]$ and a totally cyclic orientation σ of $G[S]/\mathsf{supp}(f)$.

4. $T_G(1,1) = \sum_{S \subset E} \bar{\theta}^*_{G/S}(0) \cdot \bar{\varphi}^*_{G[S]}(0)$. By Theorem 3.10.10, $\bar{\theta}^*_{G/S}(0)$ is the number of in-codegree sequences of acyclic orientations of G/S. By Theorem 3.10.7, $\bar{\varphi}^*_{G[S]}(0)$ is the number of in-degree sequences of totally cyclic orientations of $G[S]$.

In the following let $V(G/S)$ denote the vertex set of the graph G/S, $E(G[S])$ the edge set of the graph $G[S]$, etc. Also recall that $c(G[S])$ denotes the number of components of G/S while $c(G) = c$ denotes the number of components of G. If $S \subset E$, then the following identities hold:

$$c(G/S) = c(G) = c, \quad |E(G[S])| = |S|, \quad |V(G[S])| = |V|, \quad |V(G/S)| = c(G[S]).$$

5. We calculate

$$\begin{aligned} T_G(1-x, 1-y) &= \sum_{S \subset E} (-1)^{|V(G/S)| - c(G/S)} \bar{\theta}_{G/S}(x) \cdot (-1)^{|E(G[S])| - |V(G[S])| + c(G[S])} \bar{\varphi}_{G[S]}(y) \\ &= (-1)^{-c-|V|} \sum_{S \subset E} (-1)^{|S|} \bar{\theta}_{G/S}(x) \cdot \bar{\varphi}_{G[S]}(y). \end{aligned}$$

and the claim follows by the definition of $\bar{\theta}_{G/S}(x)$ and $\bar{\varphi}_{G[S]}(y)$.

6. We calculate

$$\begin{aligned} T_G(1-x, 1+y) &= \sum_{S \subset E} (-1)^{|V(G/S)| - c(G/S)} \bar{\theta}_{G/S}(x) \cdot \bar{\varphi}^*_{G[S]}(y) \\ &= (-1)^{-c} \sum_{S \subset E} (-1)^{c(G[S])} \bar{\theta}_{G/S}(x) \cdot \bar{\varphi}^*_{G[S]}(y). \end{aligned}$$

and the claim follows by the definition of $\bar{\theta}_{G/S}(x)$ and Theorem 3.3.5.

7. We calculate

$$\begin{aligned} T_G(1+x, 1-y) &= \sum_{S \subset E} \bar{\theta}^*_{G/S}(x) \cdot (-1)^{|E(G[S])| - |V(G[S])| + c(G[S])} \bar{\varphi}_{G[S]}(y) \\ &= (-1)^{-|V|} \sum_{S \subset E} (-1)^{|S| + c(G[S])} \bar{\theta}^*_{G/S}(x) \cdot \bar{\varphi}_{G[S]}(y). \end{aligned}$$

and the claim follows by the definition of $\bar{\varphi}_{G[S]}(x)$ and Theorem 3.3.6.

8. We calculate

$$\begin{aligned} T_G(1-x, 1) &= \sum_{S \subset E} (-1)^{|V(G/S)| - c(G/S)} \bar{\theta}_{G/S}(x) \cdot \bar{\varphi}^*_{G[S]}(0) \\ &= (-1)^{-c} \sum_{S \subset E} (-1)^{c(G[S])} \bar{\theta}_{G/S}(x) \cdot \bar{\varphi}^*_{G[S]}(0). \end{aligned}$$

and the claim follows by the definition of $\bar{\theta}_{G/S}(x)$ and Theorem 3.10.7.

9. We calculate

$$\begin{aligned} T_G(1, 1-y) &= \sum_{S\subset E} \bar{\theta}^*_{G/S}(0)\cdot(-1)^{|E(G[S])|-|V(G[S])|+c(G[S])}\bar{\varphi}_{G[S]}(y) \\ &= (-1)^{-|V|}\sum_{S\subset E}(-1)^{|S|+c(G[S])}\bar{\theta}^*_{G/S}(0)\cdot\bar{\varphi}_{G[S]}(y). \end{aligned}$$

and the claim follows by the definition of $\bar{\varphi}_{G/S}(x)$ and Theorem 3.10.10. □

3.12. A Combinatorial Proof of the Tutte Interpretation

In this section we give a combinatorial proof of Theorem 3.11.5, our interpretation of $T_G(1+x, 1+y)$. The approach is similar to that in Section 3.8: we show that the counting function, that we claim is identical to the Tutte polynomial, fulfills the correct deletion/contraction formula. Surprisingly, the combinatorial proof of Theorem 3.11.5 is much simpler than the combinatorial proof of Theorem 3.3.5.

If we define the sets $\mathsf{T}_G(x,y)$ by

$$\begin{aligned} \mathsf{T}_G(x,y) = \{(S,t,f) \quad : \quad & S\subset E, \\ & t \text{ a } \mathbb{Z}_x\text{-tension on } G/S, \\ & f \text{ a } \mathbb{Z}_y\text{-flow on } G[S]\}, \end{aligned}$$

for $0 < x, y \in \mathbb{N}$, then Theorem 3.11.5 simply states $T_G(1+x, 1+y) = \#\mathsf{T}_G(x,y)$ for all $0 < x, y \in \mathbb{N}$. By the deletion/contraction formula for the Tutte polynomial, all we have to show is the following:

1. If $E = \emptyset$, then $\#\mathsf{T}_G(x,y) = 1$.
2. If $e \in E$ is a bridge, then $\#\mathsf{T}_G(x,y) = (1+x)\cdot\#\mathsf{T}_{G/e}(x,y)$.
3. If $e \in E$ is a loop, then $\#\mathsf{T}_G(x,y) = (1+y)\cdot\#\mathsf{T}_{G\backslash e}(x,y)$.
4. If $e \in E$ is neither a loop nor a bridge, then $\#\mathsf{T}_G(x,y) = \#\mathsf{T}_{G/e}(x,y) + \#\mathsf{T}_{G\backslash e}(x,y)$.

For any statement A we will denote by $[A]$ the number 1 if A holds and 0 if A does not hold. Using this shorthand and the fact that if e is a loop or a bridge then $T_{G_{\backslash e}} = T_{G_{/e}}$, we can write what we have to show more compactly as

$$\#\mathsf{T}_G(x,y) = x^{[e \text{ is a bridge}]}\#\mathsf{T}_{G\backslash e}(x,y) + y^{[e \text{ is a loop}]}\#\mathsf{T}_{G/e}(x,y). \tag{3.8}$$

Before we show that this identity holds, we have to work out how the $\mathbb{Z}_k$-tensions on G and on $G_{\backslash e}$ are related, just as we did in Section 3.8 for $\mathbb{Z}_k$-flows.

Given a map $f : E \to \mathbb{Z}_k$ and a set $S \subseteq E$ we define $f|_{G_{/S}}$ and $f|_{G_{\backslash S}}$ to be the maps obtained by restricting f to the respective edge sets of $G_{/S}$ and $G_{\backslash S}$. A **fiber** of a map f is the set $f^{-1}(z)$ for any z in the image.

3.12.1. Lemma.

If t is a $\mathbb{Z}_k$-tension on G, then $t|_{G\setminus e}$ is a $\mathbb{Z}_k$-tension on $G \setminus e$. If e is a bridge, every fiber of the map $t \mapsto t|_{G\setminus e}$ has cardinality k. Otherwise every fiber of the map $t \mapsto t|_{G\setminus e}$ has cardinality 1.

Proof. A cycle in $G\setminus e$ is also a cycle in G. If $\langle C, t\rangle = 0$ holds for every cycle C of G, then it also holds for every cycle of $G\setminus e$. So $t|_{G\setminus e}$ is a $\mathbb{Z}_k$-tension on $G\setminus e$.

Now suppose e is not a bridge in G. Let t' be a $\mathbb{Z}_k$-tension on $G\setminus e$. Which $\mathbb{Z}_k$-tensions t on G have $t|_{G\setminus e} = t'$? Necessarily, $t(e') := t'(e')$ for all $e' \neq e$. All we have to show is that there is a unique choice of $t(e)$ such that t is a tension. Now as e is not a bridge, e lies on a cycle C. The weights of all other edges on C are fixed. As $\mathbb{Z}_k$ is a group, there is a unique choice of $t(e)$ such that $\langle C, t\rangle = 0$. $t(e)$ does not depend on the choice of C, as $\langle C', t\rangle = \langle C', t'\rangle = 0$ for all cycles C' that do not contain e.

If e is a bridge, then e does not lie on any cycle and so we can choose $t(e) \in \mathbb{Z}_k$ arbitrarily. □

Now the proof of our interpretation of the Tutte polynomial is easy.

Proof of Theorem 3.11.5. We have to show that (3.8) holds. To that end we define a map

$$\begin{aligned} \mathsf{T}_G(x,y) &\to \mathsf{T}_{G\setminus e}(x,y) \uplus \mathsf{T}_{G/e}(x,y) \\ (S,t,f) &\mapsto \begin{cases} (S, t|_{G\setminus e}, f) & \in \mathsf{T}_{G\setminus e}(x,y) \quad \text{if} \quad e \notin S \\ (S\setminus e, t, f|_{G/e}) & \in \mathsf{T}_{G/e}(x,y) \quad \text{if} \quad e \in S \end{cases} \end{aligned}$$

By Lemma 3.12.1 a fiber over $\mathsf{T}_{G\setminus e}(x,y)$ has cardinality x if e is a bridge and cardinality 1 otherwise. As we have seen in Section 3.8 a fiber over $\mathsf{T}_{G/e}(x,y)$ has cardinality y if e is a loop and cardinality 1 otherwise. Thus

$$\#\mathsf{T}_G(x,y) = x^{[e \text{ is a bridge}]} \#\mathsf{T}_{G\setminus e}(x,y) + y^{[e \text{ is a loop}]} \#\mathsf{T}_{G/e}(x,y)$$

for any $e \in E$. □

Chapter 4.

Counting Polynomials as Hilbert Functions

Steingrímsson [Ste01] showed that the proper $k+1$-colorings of a graph G are in bijection with the monomials of degree k in a polynomial ring $\mathbb{K}[x_1, \ldots, x_n]$ that lie inside a square-free monomial ideal I_2, but outside a square-free monomial ideal I_1. In other words, he showed that the chromatic polynomial χ_G of G is the Hilbert function of a relative Stanley-Reisner ideal. To this end, he used a clever combinatorial construction to describe the ideals I_1 and I_2 explicitly. The question we address in this chapter is this:

> Which of our five counting polynomials are Hilbert functions of relative Stanley-Reisner ideals?

Our answer is simple:

> All five are.

In particular we are able to improve Steingrímsson's result and show that $\chi_G(k)$ itself is such a Hilbert function, not only the shifted chromatic polynomial $\chi_G(k+1)$ as Steingrímsson showed.

It should be noted that Hilbert functions of relative Stanley-Reisner ideals are more general objects than Hilbert functions of standard graded algebras. The Hilbert function H_R of a standard graded algebra R necessarily has $\mathsf{H}_R(0) = 1$. This is not true of Hilbert functions of relative Stanley-Reisner ideals and this is not true of the chromatic polynomial. However, for large $k > N$ the values $\mathsf{H}_R(k)$ of the Hilbert function of standard graded algebra coincides with the values of a polynomial, which is called the Hilbert polynomial of R. So it might still be possible that there exists a standard graded algebra R whose Hilbert *polynomial* is the chromatic polynomial even though its Hilbert *function* is not.[1] Brenti [Bre98] conjectured

[1] Note that even if the Hilbert function of a relative Stanley-Reisner ideal is a polynomial, there do not in general exist standard graded algebras whose Hilbert *polynomial* coincides with the Hilbert function of the relative Stanley-Reisner ideal.

the existence of such an algebra and Almkvist gave a nonconstructive proof.[2] However, even a constructive proof of this result, may not be of as much interest as Steingrímsson's construction, even though a graded algebra is at least algebraically a more natural object than a relative Stanley-Reisner ideal. Steingrímsson [Ste01] argues: "The structure of such a graded algebra, however, will not necessarily be closely related to the colorings of G, since its monomials of degree less than the chromatic number of G cannot correspond to colorings of G." Thus, throughout this chapter we will concern ourselves with Hilbert *functions* of relative Stanley-Reisner ideals instead of Hilbert *polynomials* of graded algebras.

Back to our problem. How do we show our counting polynomials to be Hilbert functions of this type? We approach this problem from several different directions.

First, we consider only the integral tension polynomial θ_G and give an explicit combinatorial proof in the spirit of Steingrímsson's work that θ_G is a Hilbert function of this type. Our starting point is the observation that in Steingrímsson's construction the variables in $x_1, \ldots, x_n$ are in bijection with the lattice points on the boundary of the unit cube. Consequently, in the tension case we set up our polynomial ring $\mathbb{K}[x_1, \ldots, x_n]$ such that the variables are in bijection with the lattice points in tension space that lie on the boundary of the $[-1, 1]^E$ cube and develop the proof from there.

Second, we take a step back and start over by asking which Ehrhart functions are Hilbert functions of relative Stanley-Reisner ideals. Using machinery from Ehrhart theory, we quickly arrive at a sufficient criterion:

> The Ehrhart function of a relative polytopal complex in which all faces are compressed is the Hilbert function of a relative Stanley-Reisner ideal.

As was already mentioned, all five of our counting polynomials can be captured as Ehrhart functions of certain polytopal complexes - this was known for the chromatic and the integral flow and tension polynomials and we showed this for the modular flow and tension polynomials in Chapter 3. As it turns out all these polytopal complexes are of the type required by our criterion. So we can conclude that all five of our counting polynomials are Hilbert functions of Steingrímsson's type.

Third, we compare our ad hoc construction in the tension case with the argument via Ehrhart theory. It turns out that the ad hoc construction can be generalized to give another, more combinatorial proof of our sufficient criterion. The difference between the geometric and the combinatorial proof is the following. In the geometric proof we first triangulate our polytopal complex and then construct the ideal corresponding to the resulting simplicial complex. In the combinatorial proof we define an ideal corresponding to the *polytopal* complex right away and then we refine this ideal using an algebraic construction. These two constructions are related by the well-known correspondence between pulling triangulations and initial ideals of certain toric ideals with respect to the reverse lexicographic term order (see Sturmfels [Stu96]).

This combinatorial proof allows two variations. We can impose fewer requirements on our

[2] This result is attributed to Almkvist by Steingrímsson [Ste01]. The author was not able to obtain Almkvist's work himself.

polytopal complex and obtain a Hilbert function with weaker properties. If we only require our faces to be normal, we lose that the ideals are square-free. If we do not impose any constraints on our faces we also have to use a different grading of our polynomial ring.

These results are our main contribution in this chapter.

Steingrímsson [Ste01] went on to define the coloring complex of a graph to be the simplicial complex given by the square-free monomial ideal I_2. In the case of colorings the ideal I_1 has a simple description, so bounds on the f-vector of the colorings complex translate into bounds on the coefficients of the chromatic polynomial. The articles [Jon05],[Hul07],[HS08], building on Steingrímsson's work, have mainly dealt with showing various properties of the coloring complex. Steingrímsson himself gave a combinatorial description of the coloring complex and determined its Euler characteristic to be the number of acyclic orientations of G. To some extend this was already known: Welker observed that the coloring complex of a graph $G = (V, E)$ is the same as a complex appearing in the article [HRW98] by Herzog, Reiner and Welker, where this complex is shown to be homotopy equivalent to a wedge of spheres of dimension $|V| - 3$ and the number of spheres is the number of acyclic orientations of G minus one. Jonsson [Jon05] showed the coloring complex to be constructible and hence Cohen-Macaulay. This result was improved by Hultman [Hul07] who showed the coloring complex to be shellable and by Hersh and Swartz [HS08] who showed that the coloring complex has a convex ear decomposition. These results translate into bounds on the coefficients of χ_G.

The articles [HRW98], [Hul07] and [HS08] draw a connection between the coloring complex and the combinatorics of hyperplane arrangements: The coloring complex arises by intersecting a subarrangement of a certain central hyperplane arrangement with the unit sphere. In particular this relates the characteristic polynomial of the subarrangement to the chromatic polynomial. We take slightly different geometric perspective and consider lattice polytopes and Ehrhart polynomials instead of hyperplane arrangements and characteristic polynomials. A slogan describing our approach might be: Instead of intersecting the arrangement with the unit sphere we intersect it with the boundary of the unit cube.[3] However, we will not explore the relationship to the world of hyperplane arrangements. We motivate our use of Ehrhart theory directly from Steingrímsson's original construction.

In the cases of integral tensions and flows the task of obtaining bounds on the coefficients of the corresponding counting polynomials is more complicated as both the complex determined by I_2 and the complex determined by I_1 have a nontrivial structure. So what we have to say regarding the structure of these complexes and bounds on the coefficients of these counting polynomials should be regarded as only the first few steps of the investigation. Much room is left for future research.

On the one hand, we take a look at the tension complex (the complex determined by I_2 in the case of integral tensions) and the tension polytope (the complex determined by I_1 in the case of integral tensions). We show that the tension complex is homeomorphic to the coloring complex and we give a complete combinatorial characterization of the face lattice of the tension polytope. On the other hand we, give a complete characterization of Hilbert functions of relative Stanley-Reisner ideals in terms of their coefficients. This does provide

[3] Or the boundary of the $[-1, 1]^d$ cube.

new bounds on the coefficients of our five counting polynomials. However this analysis does not take the additional topological and combinatorial structure that the defining complexes have into account, so as far as bounds on the coefficients are concerned there is, as was already mentioned, much room for improvement.

The results in this chapter are joint work with Aaron Dall. The only exception is the case of non-standard gradings in Theorem 4.7.1, which is joint work with Raman Sanyal.

4.1. *Hilbert Functions*

A thorough treatment of the algebraic background of this chapter is out of scope of this thesis. As references we refer to [Eis94],[MS05] and [Sta96]. Fortunately, many of the algebraic objects we deal with have a geometric object associated with them, which make the contents of this chapter accessible also to readers without a background in commutative algebra. In this section we give brief definitions of the most important objects. In particular we define what Hilbert Functions are. We also mention a first few of the geometric correspondences. The connection between algebra and geometry will be explored more fully in Section 4.4.

Let $\mathbb{K}$ be a field. We denote by $\mathbb{K}[x] := \mathbb{K}[x_1, \ldots, x_n]$ the polynomial ring over $\mathbb{K}$ in n variables and by $\mathbb{K}[z^{\pm 1}] := \mathbb{K}[z_1^{\pm 1}, \ldots, z_d^{\pm 1}]$ the Laurent polynomial ring in d variables. Throughout the text we will write $x^a = x_1^{a_1} \cdot \ldots \cdot x_n^{a_n}$ for any $a \in \mathbb{N}^n$ and $z^u = z_1^{u_1} \cdot \ldots \cdot z_d^{u_d}$ for any $u \in \mathbb{Z}^d$. Thus the monomials in $\mathbb{K}[x]$ are in bijection with the vectors in $\mathbb{N}^n$ and the monomials in $\mathbb{K}[z^{\pm 1}]$ are in bijection with the vectors in $\mathbb{Z}^d$. Simply put, monomials correspond to lattice points. This correspondence is fundamental for this chapter.

A **monomial ideal** in $\mathbb{K}[x]$ is an ideal $I = \langle x^{a_1}, \ldots, x^{a_l} \rangle$ generated by monomials. The **dominance order** $\leq$ on $\mathbb{N}^n$ is defined componentwise: $a \leq b$ iff $a_i \leq b_i$ for all indices i. The set of monomials in a monomial ideal I, viewed as a subset of $\mathbb{N}^n$, is closed under passing to larger vectors in the dominance order. If $x^a \in I$ and $a \leq b$, then $x^b \in I$. Conversely, any set of vectors in $\mathbb{N}^n$ that is closed under passing to larger vectors in the dominance order corresponds to a monomial ideal. The generators of I correspond to the minimal elements with respect to the dominance order. As any monomial ideal is finitely generated, it follows that any subset of $\mathbb{N}^n$ has only finitely many $\leq$-minimal elements.

A monomial x^a is **square-free** if $a \leq 1$, that is if a is a $0,1$-vector. A monomial ideal is **square-free** if it is generated by square-free monomials. The generators of a square-free monomial ideal can be viewed as lattice points in the cube $[0,1]^n$. The set of lattice points in $[0,1]^n$ *not* in I is closed under passing to smaller vectors in the dominance order. Viewing these vectors as characteristic vectors of subsets of $\{1, \ldots, n\}$ we obtain a collection of sets on $\{1, \ldots, n\}$ that is closed under taking subsets. In other words, this defines an abstract simplicial complex on the ground set $\{1, \ldots, n\}$. Conversely any such abstract simplicial complex Δ defines a square-free monomial ideal I_Δ which is known as the **Stanley-Reisner ideal** of Δ. Much more will be said on this topic in Section 4.4. For now it suffices to say that the square-free monomial ideals in $\mathbb{K}[x]$ are in bijection with the abstract simplicial complexes on the ground set $\{1, \ldots, n\}$.[4]

$\mathbb{K}[x] =: S$ can be written as a direct sum of $\mathbb{K}$-vector spaces

$$S = \bigoplus_{a \in \mathbb{N}^n} S_a$$

where each S_a is the $\mathbb{K}$-linear subspace generated by the monomial x^a. Since the product of any two monomials $x^a \in S_a$ and $x^b \in S_b$ is $x^a x^b = x^{a+b} \in S_{a+b}$, we have $S_a \cdot S_b \subset S_{a+b}$ for all $a, b \in \mathbb{N}^n$. A ring with such a decomposition, we call $\mathbb{N}^n$**-graded** and the decomposition itself we call a $\mathbb{N}^n$**-grading**.

Let M be an S-module, where S is $\mathbb{N}^n$-graded. A $\mathbb{N}^n$**-grading** of M is a decomposition of M as a direct sum of abelian groups

$$M = \bigoplus_{a \in \mathbb{N}^n} M_a$$

such that $S_a M_b \subset M_{a+b}$ for all $a, b \in \mathbb{N}^n$. M along with a given $\mathbb{N}^n$-grading we call $\mathbb{N}^n$**-graded**. Note that since $S = \mathbb{K}[x]$, any summand M_a has the structure of a $\mathbb{K}$-vector space. The dimension of this space we denote by $\mathsf{dim}_{\mathbb{K}} M_a$. The **fine Hilbert function** is the function $\mathsf{H}_M : \mathbb{N}^n \to \mathbb{N}$ that maps $a \mapsto \mathsf{dim}_{\mathbb{K}} M_a$.

As an example consider any monomial ideal I. I can be written as a direct sum $I = \bigoplus_{a \in \mathbb{N}^n} I_a$ where I_a is the 1-dimension $\mathbb{K}$-linear subspace generated by x^a provided $x^a \in I$. If $x^a \notin I$, then $I_a = \{0\}$ is the 0-dimensional $\mathbb{K}$-linear subspace. $S_a I_b \subset I_{a+b}$, so I is an $\mathbb{N}^n$-graded module. Its fine Hilbert function is just the characteristic function of the set of lattice points in I.

A second example is the quotient $\mathbb{K}[x]/I$. This quotient inherits its grading from $\mathbb{K}[x]$, that is $\mathbb{K}[x]/I = \bigoplus_{a \in \mathbb{N}^n} S_a/I_a$ where S_a/I_a is a quotient of $\mathbb{K}$-vector spaces. Unless otherwise stated, quotients $\mathbb{K}[x]/I$ will always inherit the $\mathbb{N}^n$-grading of $\mathbb{K}[x]$ in this fashion. The fine Hilbert function $\mathsf{H}_{\mathbb{K}[x]/I}$ of $\mathbb{K}[x]/I$ is given by $\mathsf{H}_{\mathbb{K}[x]/I}(a) = 1$ if $x^a \in I_a$ and $\mathsf{H}_{\mathbb{K}[x]/I}(a) = 0$ if $x^a \notin I_a$. Viewed as a lattice point set $\mathbb{K}[x]/I$ is the complement of I. This set is closed under passing to vectors smaller in the dominance order.

The ring S also has a grading by degree. The degree of a monomial $x^a \in \mathbb{K}[a]$ is defined as usual as $\deg(x^a) := \sum_{i=1}^n a_i$. Since $a \in \mathbb{N}^n$ we can also write this as $\deg(x^a) = ||a||_1$. Then

$$S = \bigoplus_{k \in \mathbb{N}} S_k \quad \text{where} \quad S_k = \bigoplus_{\substack{a \in \mathbb{N}^n \\ ||a||_1 = k}} S_a \;\; \text{for all } k \in \mathbb{N}.$$

Now $\deg(x^a x^b) = \deg(x^a) + \deg(x^b)$ so $S_i S_j \subset S_{i+j}$ for all $i, j \in \mathbb{N}$. A ring with such a decomposition we call $\mathbb{N}$**-graded** and this grading of S we call the **standard grading**.

Given the standard grading of S, we can proceed as above to define the $\mathbb{N}$-grading of an S-module M. A decomposition of M as a direct sum $M = \bigoplus_{k \in \mathbb{N}} M_k$ such that $S_i M_j \subset M_{i+j}$

[4] This is not quite true. If we disallow the empty simplicial complex, but allow the simplicial complex that consists only of the empty set, then we cannot represent the ideal $\mathbb{K}[x] = \langle 1 \rangle$ but we can represent $\langle x_1, \ldots, x_n \rangle$. This fine point will not be of importance though.

is called an $\mathbb{N}$**-grading** of M. An S-module M is called $\mathbb{N}$**-graded** if it is given along with a fixed $\mathbb{N}$-grading. The **Hilbert function** of M is the function $\mathsf{H}_M : \mathbb{N} \to \mathbb{N}$ given by

$$\mathsf{H}_M(k) = \dim_{\mathbb{K}} M_k.$$

These are the functions we will be dealing with in the rest of this chapter. The guiding question will be: Which counting functions are Hilbert functions? Although, to be precise, we will be looking for Hilbert functions that arise from a particular kind of module. We will say more about this in the next section, where we present Steingrímsson's result, which captures the chromatic polynomial of a graph as a Hilbert function.

We conclude this section by considering, again, our two examples. First, let I be a monomial ideal. Unless otherwise stated, any monomial ideal is understood to be equipped with the standard grading $I = \bigoplus_{k\in\mathbb{N}} I_k$ where I_k is the $\mathbb{K}$-vector space spanned by all monomials of degree k that appear in I. If there are no monomials of degree k in I, then I_k is 0-dimensional. The Hilbert function H_M of I maps $k \in \mathbb{N}$ to the number of monomials in I that are of degree k. Viewing I as a set of lattice points, this is just the number of lattice points in I that are contained in the simplex $\{x \in \mathbb{R}^n | x \geq 0, \sum_{i=1}^n x_i = k\}$.

Second, let $\mathbb{K}[x]/I$ be the quotient of the polynomial ring by a monomial ideal. This quotient is understood to inherit the $\mathbb{N}$-grading of $\mathbb{K}[x]$ and I via $\mathbb{K}[k]/I = \bigoplus_{k\in\mathbb{N}} S_k/I_k$ where, again, S_k/I_k is a quotient of $\mathbb{K}$-vector spaces. S_k/I_k is generated by all monomials x^a of degree k with $x^a \notin I_k$. So the Hilbert function $\mathsf{H}_{\mathbb{K}[x]/I}$ of $\mathbb{K}[x]/I$ maps $k \in \mathbb{N}$ to the number of monomials of degree k not in I, which is just the number of lattice points in the simplex $\{x \in \mathbb{R}^n | x \geq 0, \sum_{i=1}^n x_i = k\}$ that are *not* contained in I.

4.2. Steingrímsson's Construction

Steingrímsson [Ste01] showed that for any graph G the chromatic polynomial $\chi_G(k+1)$ shifted by one is the Hilbert function of a module with a particular structure.

4.2.1. Theorem. (Steingrímsson [Ste01, Theorem 9])

For any graph G, there exists a number n, a square-free monomial ideal I_1 in the polynomial ring over n variables $\mathbb{K}[x] = \mathbb{K}[x_1, \ldots, x_n]$ and a square-free monomial ideal I_2 in $\mathbb{K}[X]/I_1$ such that

$$\mathsf{H}_{I_2}(k) = \chi_G(k+1)$$

for all $k \in \mathbb{N}$, where H_{I_2} denotes the Hilbert function of I_2 with respect to the standard grading and χ_G denotes the chromatic polynomial of G.

In other words Steingrímsson constructed two square-free monomial ideals $I_2 \supset I_1$ in $\mathbb{K}[x]$ such that the monomials of degree k in I_2 but not in I_1 are in bijection with the proper $k+1$-colorings of G. Here we note that a monomial ideal $I_2 \supset I_1$ in $\mathbb{K}[x]$ can be viewed as a monomial ideal I_2' in $\mathbb{K}[x]/I_1$. $I_2' \subset \mathbb{K}[x]/I_1$ is simply the image of I_2 under the projection

$\mathbb{K}[x] \to \mathbb{K}[x]/I_1$. Conversely, a monomial ideal $I_2' \subset \mathbb{K}[x]/I_1$ can be viewed as a monomial ideal $I_2 \subset \mathbb{K}[x]$ with $I_2 \supset I_1$. I_2' can be written as a direct sum $I_2' = \bigoplus_{a \in \mathbb{N}^n} I_{2,a}'$ where $I_{2,a}'$ is a $\mathbb{K}$-linear subspace of $S_a/I_{1,a}$. We recall that $S_a/I_{1,a}$ can either be 0- or 1-dimensional, so $I_{2,a}'$ can either be 0- or 1-dimensional and $I_{2,a}'$ has to be 0-dimensional if $S_a/I_{1,a}$ is. So $I_2 = \langle x^a | \dim I_{2,a}' = \dim S_a/I_{1,a} \rangle$ is the unique monomial ideal in $\mathbb{K}[x]$ with $I_2 \supset I_1$ whose image under the projection $\mathbb{K}[x] \to \mathbb{K}[x]/I_1$ is I_2'. From now on we will no longer distinguish between I_2 and I_2' and regard I_2 either as a monomial ideal in $\mathbb{K}[x]$ with $I_2 \supset I_1$ or as a monomial ideal in $\mathbb{K}[x]/I_1$, whichever is more convenient.

Phrased more geometrically, Steingrímsson constructed two sets of lattice points $I_2 \supset I_1$ in $\mathbb{N}^n$, both closed under passing to larger vectors in the dominance order and both with the property that their minimal elements are in $\{0,1\}^n$, such that the set $I_2 \setminus I_1$ intersected with the simplex $\{x | x \geq 0, \sum_{i=1}^n x_i = k\}$ is in bijection with the proper $k+1$-colorings of G.

To achieve this, he proceeded as follows.[5]

Proof. Let $G = (V, E)$ be a graph with vertices labeled $1, \ldots, m$. Let $\mathbb{K}[x]$ denote the polynomial ring with 2^m variables that are indexed with the subsets of $\{1, \ldots, m\} =: [m]$. So the variables x_S in $\mathbb{K}[x]$ stand for sets $S \subset [m]$. Whenever we refer to as a variable x_S as a "set" we will mean the set S. Now let

$$I_1 = \langle x_S x_T | S \not\subset T \text{ and } S \not\supset T \rangle$$

and consider $\mathbb{K}[x]/I_1$. Any monomial x^a that is not zero in $\mathbb{K}[x]/I_1$ will be a product of variables whose corresponding sets are pairwise comparable. This means that the sets corresponding to the variables appearing in x^a form a chain. Every monomial x^a that is not zero in $\mathbb{K}[x]/I_1$ can be written as

$$x^a = x_{S_1}^{a_{S_1}} \cdot x_{S_2}^{a_{S_2}} \cdot \ldots \cdot x_{S_l}^{a_{S_l}} \tag{4.1}$$

for some l where $a_{S_i} > 0$ for $1 \leq i \leq l$ and

$$[m] \supset S_1 \supsetneq S_2 \supsetneq \ldots \supsetneq S_l \supset \emptyset.$$

The monomial $x^0 = 1$ is seen as corresponding to the empty chain $[m] \supset \emptyset$ with $l = 0$.

We will associate each such monomial of degree k with a $k+1$-coloring of G. The underlying idea is simple and quite clever, but it takes a moment to digest. We give two equivalent definitions of the coloring c corresponding to x^a.

1. Algorithmically c is determined as follows. Each vertex starts out with color 0. We process (4.1) from left to right. For every monomial $x_{S_i}^{a_{S_i}}$ we encounter, we increase the color of all vertices in S_i by a_{S_i}.

2. Put $S_0 := \emptyset$ and $S_{l+1} := V$. Every vertex is contained in a unique set $S_i \setminus S_{i+1}$ for $0 \leq i \leq l$. If $v \in S_i \setminus S_{i+1}$, then $c(v) = \sum_{j=1}^i a_{S_j}$.

[5] We give Steingrímsson's construction with two small modifications. First, we label our k+1 colors 0,...,k. And second, we consider our chains $S_i \supsetneq S_{i+1}$ to be descending, whereas he considers the chains to be ascending. Conceptually, these changes do not make any difference.

So, loosely speaking, the color of a vertex v is the number of factors x_S with $v \in S$ that appear in the monomial. Note that the definition associates the monomial $x^0 = 1$ with the constant coloring in which every vertex receives the color 0.

Clearly, every coloring defined this way uses at most the colors 0 through $\deg(x^a)$. If $S_1 = [m]$, then the color 0 is not used and the smallest color that is used is a_{S_1}. If $S_l = \emptyset$, then the color $\deg(x^a)$ is not used and the largest color that is used is $\deg(x^a) - a_{S_l}$.

So we have seen that every non-zero monomial in $\mathbb{K}[x]/I_1$ determines a coloring. Conversely a $k+1$-coloring c gives rise to such a monomial in the following way. For any color $i \in \{0,\ldots,k\}$ let $c^{-1}(i)$ denote the set of vertices with color i and let $C_{\geq i} = \bigcup_{j=i}^{k} c^{-1}(j)$ denote the set of vertices that have color at least i. Now we define

$$x^a := x_{C_{\geq 1}} \cdot x_{C_{\geq 2}} \cdot \ldots \cdot x_{C_{\geq k}}.$$

Note that sets $C_{\geq i}$ and $C_{\geq j}$ may coincide, which leads to variables with higher exponents in x^a. By construction $\deg(x^a) = k$ and the coloring corresponding to this monomial is c.

Let us now turn to the question which monomials define proper colorings. Let c be the coloring determined by the monomial from (4.1). Two vertices v and w have the same color under c if and only if there exists an $0 \leq i \leq l$ such that $v, w \in S_i \setminus S_{i+1}$. So c is proper if and only if for all $0 \leq i \leq l$ the induced graph on vertex set $S_i \setminus S_{i+1}$ has no edges, that is, if the vertex set $S_i \setminus S_{i+1}$ is independent. Monomials x^a, non-zero in $\mathbb{K}[x]/I_1$, such that the sets $S_i \setminus S_{i+1}$ are all independent are called **coloring monomials**. Now comes the crux of the argument:

4.2.2. If x^a is a coloring monomial and x^b any monomial such that $x^a x^b \notin I_1$, then $x^a x^b$ is a coloring monomial.

The reason is simple: If we multiply x^a with x^b the sets $S_i \setminus S_{i+1}$ can only get smaller. So the ideal

$$I_2 = \langle x^a | x^a \text{ is a coloring monomial} \rangle$$

contains *precisely* the coloring monomials.

We have thus constructed a monomial ideal $I_1 \subset \mathbb{K}[x]$ and a monomial ideal $I_2 \subset \mathbb{K}[x]/I_1$ such that the monomials $x^a \in I_2$ of degree k that are non-zero in $\mathbb{K}[x]/I_1$ are in bijection with the $k+1$-colorings of G. Clearly, both ideals are square free, as whether a monomial x^a is contained in either depends only on which variables appear in x^a and not on their multiplicity. □

It is crucial to note one thing about the above construction: The variables x_S correspond to 0,1-vectors in $[0,1]^{|V|}$, namely the characteristic vectors of the sets S. Let z_S, defined by $z_S(v) = 1$ if $v \in S$ and $z_S(v) = 0$ otherwise, denote the characteristic vector of a set S. If S ranges over all subsets of V, then z_S ranges over all vertices of the cube $[0,1]^V$. The monomial x^a given in (4.1) corresponds to the coloring

$$c = a_{S_1} z_{S_1} + a_{S_2} z_{S_2} + \ldots + a_{S_l} z_{S_l},$$

and thus the coloring monomials are merely integral representations of all colorings. The ideal I_1 is chosen in such a way that every coloring has a unique integral representation and all representations have the right degree, while the ideal I_2 is chosen such that only proper colorings appear. In the next sections we are going to expand upon this point of view. In Section 4.3 we will apply this observation directly to construct ideals I_1 and I_2 that capture the tension polynomial. In Section 4.4 we are going to take a step back and approach the problem from a more general and abstract point of view, identifying the Hilbert functions of the type Steingrímsson considered as Ehrhart functions of certain complexes. This turns the question whether some counting function f is a Hilbert function into the question whether there is a certain complex whose Ehrhart function is f.

4.3. The Tension Polynomial as a Hilbert Function (Combinatorially)

In this section we will plunge right in and see if we can build on the observations in Section 4.2 to obtain the integral tension polynomial as a Hilbert function. In Section 4.4 take a different and more geometric approach to this task and develop a general framework for solving this kind of problem.

We want to obtain an analogue of Theorem 4.2.1 for the integral tension polynomial as introduced in Sections 3.2 and 3.4. Our overall approach is going to be this:

> Given a graph $G = (V, E)$ we want to construct square-free monomial ideals I_1 and I_2 in some polynomial ring $\mathbb{K}[x]$ such that the monomials in $I_2 \setminus I_1$ are in bijection with the nowhere zero $\mathbb{Z}$-tensions on G. Moreover this bijection is supposed to be such that it induces a bijection between the monomials of degree k and the $\mathbb{Z}$-tensions t with $\max_{e\in E} |t(e)| = ||t||_\infty = k$.

Then the number of the monomials of degree *at most* k is the number of nowhere zero $(k+1)$-tensions on G and we are done by the following observation.

4.3.1. Let $I_1 \subset I_2$ be monomial ideals in $\mathbb{K}[x_1, \ldots, x_n]$, and let $f(k)$ be the number of monomials in $I_2 \setminus I_1$ of degree k. Let A_1 and A_2 be sets of monomials with $I_1 = \langle A_1 \rangle$ and $I_2 = \langle A_2 \rangle$. Now consider A_1 and A_2 as sets of monomials in $\mathbb{K}[x_1, \ldots, x_n, x_{n+1}]$ and define $I_1' = \langle A_1 \rangle$ and $I_2' = \langle A_2 \rangle$ in this larger ring. Then the number of monomials in $I_2' \setminus I_1'$ is $\sum_{i=0}^{k} f(i)$.

Proof. The crucial observation is that $x_1^{a_1} \cdot \ldots \cdot x_n^{a_n} \in I_1$ if and only if $x_1^{a_1} \cdot \ldots \cdot x_n^{a_n} x_{n+1}^{a_{n+1}} \in I_1'$ for any choice of a_{n+1} and the same holds for I_2. Thus

$$\begin{aligned} & \#\{x_1^{a_1} \cdot \ldots \cdot x_n^{a_n} x_{n+1}^{a_{n+1}} \in I_2' \setminus I_1' \text{ of degree } k\} \\ = & \sum_{i=0}^{k} \#\{x_1^{a_1} \cdot \ldots \cdot x_n^{a_n} \in I_2 \setminus I_1 \text{ of degree } i\} \end{aligned}$$

which completes the proof. □

Definition of the Polynomial Ring and the Bijection. What is our basic setup of the polynomial ring $\mathbb{K}[x]$ and the desired bijection? In Section 4.2 we observed that the variables x_S of the polynomial ring $\mathbb{K}[x]$ Steingrímsson used corresponded to all the $0,1$-colorings of the graph G. So to capture the integral tension polynomial, let $\mathcal{U}$ denote the set of all $0,\pm 1$-tensions on G, excluding the all zero tension. As we have already seen, the set $\mathcal{U}$ includes all the rows of the incidence matrix of G, but $\mathcal{U}$ contains more vectors than these. $\mathcal{U}$ also includes the sign vectors of all directed cuts in G, but $\mathcal{U}$ contains still more vectors than these. Now we consider the polynomial ring $\mathbb{K}[x] := \mathbb{K}[x_a : u \in \mathcal{U}]$ where we have one variable x_u for every $0,\pm 1$-tension $u \in \mathcal{U}$. This is the polynomial ring we want to construct our two ideals I_1 and I_2 in. On the other hand we view $\mathbb{Z}$-tensions as vectors in $\mathbb{Z}^E$ and this set we can identify with the Laurent polynomial ring $\mathbb{K}[z^{\pm 1}] := \mathbb{K}[z_e^{\pm 1} : e \in E]$ where we have one variable z_e for every edge $e \in E$ of the graph. The correspondence between the variable x_u in $\mathbb{K}[x]$ and the $0,\pm 1$-tensions $u \in \mathcal{U}$ now translates, on the one hand, into a linear map $\pi : \mathbb{N}^{\mathcal{U}} \to \mathbb{Z}^E$ and, on the other hand, into a homomorphism $\hat{\pi} : \mathbb{K}[x] \to \mathbb{K}[z^{\pm 1}]$, that are defined by

$$\pi : a \quad \mapsto \quad Ua$$

where U is the matrix whose columns are the elements of $\mathcal{U}$ and

$$\hat{\pi} : x^a \quad \mapsto \quad z^{Ua}$$

for all $u \in \mathcal{U}$. We define $n := |\mathcal{U}|$ and $d := |E|$ and identify $\mathbb{N}^n$ with $\mathbb{N}^{\mathcal{U}}$ and $\mathbb{Z}^d$ with $\mathbb{Z}^E$. We will call a a **representation** of Ua, x^a a representation of $\hat{\pi}(x^a)$, x^a a representation of Ua and so forth. The good news is that with this basic setup every monomial in $\mathbb{K}[x]$ represents a $\mathbb{Z}$-tension, simply because the $\mathbb{Z}$-tensions are the lattice points in tension space. Also, every $\mathbb{Z}$-tension is representable, as we already know from Section 3.4.

Now we need to construct our ideals I_1 and I_2 such that all of the following issues are taken care of.

1. The tensions corresponding to monomials $x^a \in I_2 \setminus I_1$ are nowhere zero.
2. The degrees of the monomials line up correctly.
3. $\hat{\pi}|_{I_2 \setminus I_1}$ is surjective.
4. $\hat{\pi}|_{I_2 \setminus I_1}$ is injective.
5. I_1 and I_2 are square free.

This is a lot to handle. Instead of simply giving the right definitions for I_1 and I_2 and then proving that all constraints are satisfied, we will try to motivate where the definitions come from. We will go through several candidate sets I_1^i and I_2^j, refining the definitions in each iteration to satisfy a new constraint.

Before we start with the construction, let us get a feel for the structure of sets of the form $I_2 \setminus I_1$. The monomials $x^a \in I_2 \setminus I_1$ we call legal representations or simply **legal**.

What sets of monomials are of the form $I_2 \setminus I_1$? What does it mean that a monomial x^a is in I_2 but not in I_1? What kinds of constraints on x^a can we formulate by means of these two ideals?

First of all, we recall that because I_1 and I_2 are square-free, any constraint on x^a can only depend on *which* variables appear in x^a, not on *the exponent* of that variable. In other words the constraint may only depend on $\mathsf{supp}(a)$ not on a itself.

Second, we observe that I_1 allows us to forbid certain products of variables to appear in x^a. We can formulate a constraint like "x_u and x_v may not *both* appear in x^a" by including $x_u x_v$ in I_1. On the level of exponent vectors, this means that for any $0,1$-vector b we can require of x^a that $\mathsf{supp}(b) \not\subset \mathsf{supp}(a)$.

Third, I_2 allows us to specify that at least one of a given list of products of variables has to appear in x^a. We can formulate constraints like "at least one of $x_u x_v$ or $x_u x_s x_t$ or x_w or ... must appear in x^a" by letting $I_2 = \langle x_u x_v, x_u x_s x_t, x_w, \ldots \rangle$. On the level of exponent vectors, this means that for a list $b_1, \ldots, b_l$ of $0,1$-vectors we can require of x^a that there exists an i such that $\mathsf{supp}(b_i) \subset \mathsf{supp}(a)$.

Nowhere Zero Tensions. So what does this mean in relation to our nowhere zero $\mathbb{Z}$-tensions? We have to construct I_1 and I_2 such that any legal x^a represents a nowhere zero $\mathbb{Z}$-tension. The most straightforward thing we could do is define I_2 to be the ideal *generated* by those monomials x^a that represent a nowhere zero $\mathbb{Z}$-tension.

$$I_2^1 := \langle x^a | a \in \mathbb{N}^n,\ Ua \text{ is nowhere zero} \rangle \tag{4.2}$$

By definition, this is an ideal that contains all monomials that represent a nowhere zero tensions. The problem, however, is that this ideal also contains monomials that represent tensions that **are** zero on some edge e. Suppose both $x^a, x^b \in I_2^1$ represent, respectively, nowhere zero $\mathbb{Z}$-tensions s and t. Suppose further that there is an edge e such that s_e and t_e have opposite sign. For simplicity, let us assume that $s_e = -t_e \neq 0$. Then x^{a+b} represents the $\mathbb{Z}$-tension $s+t$ which is zero on e. Both x^a and x^b are in I_2^1, as they correspond to a nowhere zero tension, and $x^{a+b} = x^a x^b \in I_2^1$, as I_2^1 is an ideal, but the corresponding tension is zero somewhere. So we have to exclude x^{a+b} using I_1: we forbid the multiplication of x^a and x^b by letting $x^{a+b} \in I_1$. This motivates the following definition.

$$I_1^1 := \langle x_u x_v \mid u, v \in \mathcal{U},\ \exists e \in E:\ u_e = -v_e \neq 0 \rangle \tag{4.3}$$

Recall that two vectors $u, v \in \mathbb{Z}^d$ are sign-compatible if there does not exist an $i \in [d]$ such that $\mathsf{sgn}(u_i) = -\mathsf{sgn}(v_i) \neq 0$. It is important to observe that the monomials $x^a \in I_1^1$ are precisely those monomials such that for any two factors x_u and x_v of x^a the vectors u and v are sign-compatible. Moreover I_1^1 is clearly square-free. To obtain a square-free I_2, we have to modify the definition of I_2^1 slightly. We put

$$I_2^2 := \langle x^a | a \in \mathbb{N}^n,\ x^a \notin I_1^1,\ Ua \text{ is nowhere zero} \rangle.$$

Clearly this change did not affect the set $I_2 \setminus I_1$, but now we can argue that I_2^2 is square-free. Let $x^a \in I_2^2$ and suppose $a_u \geq 2$. Consider the vector $a' := a - e_u$ where the u-th entry of a

has been decremented by 1. As for all $e \in \mathsf{supp}(u)$ we have $(Ua)_e \geq 2$ *by definition of* I_1, we know $\mathsf{supp}(Ua') = \mathsf{supp}(Ua)$ and so $x^{a'}$ also represents a nowhere zero tension and hence is included in I_2^2, showing that x^a is not a minimal generator of I_2^2.

The above arguments imply:

4.3.2. Ua is nowhere zero for every monomial in $x^a \in I_2^2 \setminus I_1^1$.

Handling the Degrees. As of now, the degrees do not line up as desired. If x^a and x^b represent tensions t_1 and t_2 with disjoint non-empty support, then $x^a x^b$ will represent the tension $t_1 + t_2$ which has maximal edge weight

$$||t_1 + t_2||_\infty = \max\{||t_1||_\infty, ||t_2||_\infty\} \neq ||t_1||_\infty + ||t_2||_\infty = \deg x^a x^b.$$

So we have to make sure that there is an edge e such that both $|t_1(e)| = ||t_1||_\infty$ and $|t_2(e)| = ||t_2||_\infty$. Thus we define $\mathsf{Max}(t) = \{e \in E : |t(e)| = ||t||_\infty\}$ and require that all the tensions t_i appearing in a representation have an element of $\mathsf{Max}(t_i)$ in common. We say tensions $t_1, \dots, t_l$ **stack** if the t_i are pairwise sign-compatible and $\bigcap_{i=1}^{l} \mathsf{Max}(t_i) \neq \emptyset$. Similarly we say of variables $x_{u_1}, \dots, x_{u_l}$ that they stack if the $0, \pm 1$-tensions $u_1, \dots, u_l$ stack. We say a vector a is **stacking** if the tensions in $\mathsf{supp}(a)$ stack and call a a **stacking representation** of Ua. As a stacking representation a is a combination of sign-compatible $0, \pm 1$-tensions the edges in $\mathsf{Max}(Ua)$ are precisely those that appear in the supports of all $u \in \mathsf{supp}(a)$.

4.3.3. If a is a stacking representation of Ua, then $\mathsf{Max}(Ua) = \bigcap_{u \in \mathsf{supp}(a)} \mathsf{supp}(u)$.

Our next candidate for I_1 is now

$$I_1^2 \quad := \quad \langle x^a \mid a \text{ is not a stacking representation } \rangle.$$

Recall that I_1^1 was generated by all monomials x^a such that the variables appearing in I_1^1 are not pairwise sign-compatible, so I_1^1 is contained in I_1^2. Note that whether or not the variables appearing in x^a stack depends only on $\mathsf{supp}(a)$. Thus I_1^2 is a square-free monomial ideal and for any monomial $x^a \in \mathbb{K}[x]$ we have that $x^a \notin I_1^2$ if and only if the variables appearing in x^a stack (and, consequently, are pairwise sign-compatible).

The benefit of this definition is this.

4.3.4. $\deg(x^a) = ||Ua||_\infty$ for any non-zero monomial $x^a \in \mathbb{K}[x]/I_1^2$.

Note that this change of I_1 preserves 4.3.2.

Surjectivity. Our definitions are still not final. Yet, the time has come to show that our proposed map $\hat{\pi}|_{I_2 \setminus I_1}$ is surjective: We claim that every nowhere zero $\mathbb{Z}$-tension is represented by some monomial in $I_2^2 \setminus I_1^2$. All we have to show is that every $\mathbb{Z}$-flow has a stacking representation. We are going to show a stronger statement, that gives us more control over the representation. This better control will be crucial when we show that I_1 is square-free at the end of this section.

4.3.5. Let t be a $\mathbb{Z}$-tension.

1. There is a $0, \pm 1$-tension u that is sign-compatible with t such that

$$\mathsf{Max}(t) \subset \mathsf{supp}(u) \subset \mathsf{supp}(t),$$

2. t has a stacking representation a with $a_u \geq 1$, for any such u.

This can be shown using the total unimodularity of a certain matrix and we will elaborate on this approach in Sections 4.4 and 4.5. Here, however, we give a direct combinatorial proof of this statement. It is short, constructive and has the nice property that we can illustrate it with a picture. In the opinion of the author this proof is one example showing that it is worthwhile to consider a combinatorial problem in terms of the concrete combinatorial objects, even if there is a rich and beautiful theory available that can solve the problem for you.

Proof. *1.* Let t be a $\mathbb{Z}$-tension and c any corresponding coloring. Let $w = \max\{|t(e)| : e \in E\}$ denote the absolute value of the largest (absolute) edge weight appearing in t. Now define a coloring c' by $c'(v) = c(v) \text{ div } w$ for all $v \in V$ and let u be the corresponding tension.

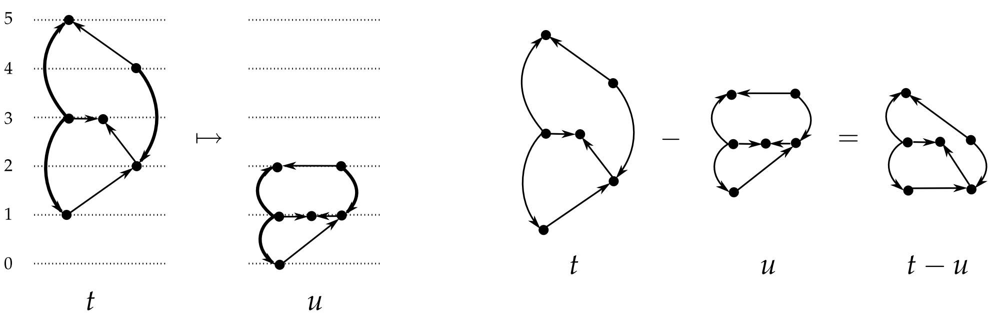

u is a $0, \pm 1$-tension by choice of w and c' with $\mathsf{Max}(t) \subset \mathsf{supp}(u)$. Moreover if $c(v_1) \leq c(v_2)$ then $c'(v_1) \leq c'(v_2)$ and thus t and u are sign-compatible. On the other hand if $c(v_1) = c(v_2)$ then $c'(v_1) = c'(v_2)$ and so $\mathsf{supp}(t) \supset \mathsf{supp}(u)$.

2. We do induction over $||t||_\infty$. If $||t||_\infty = 1$, then $u = t$ and we are done. Otherwise $t - u$ is a $\mathbb{Z}$-tension with $||t - u||_\infty = ||t||_\infty - 1$, that is sign-compatible with t. Moreover, because $\mathsf{Max}(u) \supset \mathsf{Max}(t)$ and u is a sign-compatible $0, \pm 1$-tension we know $\mathsf{Max}(t) \subset \mathsf{Max}(t - u)$.

Using 1. and the induction hypothesis, we find a stacking representation a' of $t - u$. Then $a' + e_u$ is a stacking representation of t with $u \in \mathsf{supp}(a' + e_u)$, as desired. □

Injectivity. We have now seen that the map $\hat{\pi}$ maps the set of monomials in $I_2 \setminus I_1$ surjectively onto the nowhere zero $\mathbb{Z}$-tensions. The bad news is that this map is not injective. There are $\mathbb{Z}$-tensions that have several stacking representations. An example is given in the following picture.

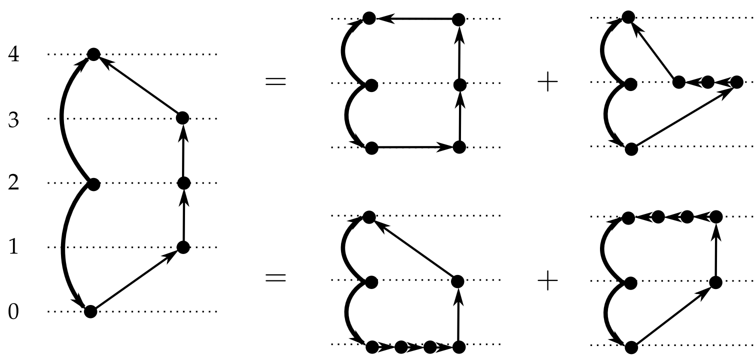

So we have to make the set $I_2 \setminus I_1$ smaller, somehow, to obtain a bijection. At the same time we have to maintain that I_2 and I_1 are square-free monomial ideals. At first glance this may seem hopeless. But this is where algebra comes to the rescue. Of all the representations of a given tension we can simply pick one and throw all the others into I_1. If we make these choices "consistently" we keep a square-free monomial ideal. More about this method is mentioned in Section 4.6 and, especially, in Sturmfels' book [Stu96].

We fix a linear order on the variables of $\mathbb{K}[x]$. A **term order** is a total order $\prec$ on the monomials in $\mathbb{K}[x]$ such that $1 \preceq x^a$ for every monomial x^a and such that $x^a \prec x^b$ implies $x^a x^c \prec x^b x^c$ for any monomials x^a, x^b, x^c. If $p(x) = \sum_{a \in S} c_a x^a$ is any polynomial, then the **initial term** of p with respect to a given term order $\prec$ is $\mathsf{in}_\prec(p) = c_{a_0} x^{a_0}$ where x^{a_0} is $\prec$-maximal among all monomials appearing in p, that is all monomials x^a with $a \in S$ and $c_a \neq 0$. For any ideal I in $\mathbb{K}[x]$ the **initial ideal** $\mathsf{in}_\prec(I) = \langle \mathsf{in}_\prec(p) | p \in I \rangle$ is the ideal generated by the initial terms of all the polynomials in I. A term order also induces a total order on $\mathbb{N}^n$ via $a \prec b$ iff $x^a \prec x^b$.

The reverse lexicographic term order is going to be of particular interest to us. To construct this term order, we first fix an arbitrary total order on the variables of $\mathbb{K}[x]$. We already used a total order on the variables implicitly when we identified monomials x^a with vectors $a \in \mathbb{N}^n$ and thought of the entries of a as being ordered from left to right or top to bottom. Our constructions are not going to depend on the chosen total order on the variables; we just pick an arbitrary one and fix it for the rest of this chapter. Also we are not going to denote this order by a symbol. Instead we are going to use phrases like "x_u comes before x_v", "x_u is left of x_v" or "x_u is above x_v". Now we define the **reverse lexicographic term order** $\prec_{\mathsf{revlex}}$ as follows. Let x^a and x^b be two monomials in $\mathbb{K}[x]$. If $\deg(x^a) < \deg(x^b)$, then $x^a \prec_{\mathsf{revlex}} x^b$. If $\deg(x^a) = \deg(x^b)$, then $x^a \prec_{\mathsf{revlex}} x^b$ if in the rightmost entry where a and b differ, a is bigger. $\prec_{\mathsf{revlex}}$ is indeed term order. Note that if we did not classify by degree first, we would not obtain a term order as 1 would not be the minimal element. We also use $\prec_{\mathsf{revlex}}$ to denote the induced total order on $\mathbb{N}^n$.

Recall that the set of non-zero monomials in $\mathbb{K}[x]/I_1^2$ gave us a set of exponent vectors in $\mathbb{N}^n$ that is mapped *surjectively* onto the $\mathbb{Z}$-tensions by U. Our goal is now to construct a subset that is mapped *bijectively* onto the $\mathbb{Z}$-tensions by implementing the following rule.

> If a $\mathbb{Z}$-tension t has several representations, pick the one that is smallest with respect to $\prec_{\mathsf{revlex}}$.

To this end we construct an ideal I_1' as follows

$$\begin{aligned} I_1' &:= \langle x^b | x^b \notin I_1^2, \exists x^a \notin I_1^2 : Ua = Ub \wedge x^a \prec_{\mathsf{revlex}} x^b \rangle \\ &= \mathsf{in}_{\prec_{\mathsf{revlex}}} \langle x^a - x^b | x^a, x^b \notin I_1^2, Ua = Ub \rangle \end{aligned}$$

and define our final version of the ideal I_1 to be

$$I_1 := I_1^2 + I_1'.$$

Note that $I_1 = I_1^2 + I_1'$ is the monomial ideal generated by the union of the generators of I_1 and I_1'. We now consider the ring $\mathbb{K}[x]/I_1$.

First of all we note that a monomial $x^b \notin I_1^2$ is in I_1' if and only if there exists a $x^a \notin I_1^2$ such that $Ua = Ub$ and $x^a \prec_{\mathsf{revlex}} x^b$. By definition, all x^b with this property are in I_1'. To see the converse, we observe that if $x^b \notin I_1^2$ and $x^a \notin I_1^2$ is such that $Ua = Ub$ and $x^a \prec_{\mathsf{revlex}} x^b$, then for any $x^c \notin I_1^2$ with $b \leq c$ we have that c is also not a $\prec_{\mathsf{revlex}}$-minimal representation with $x^c \notin I_1^2$. Let $d = c - b$. Then $Uc = U(b + d) = U(a + d)$ by linearity of U and $x^{a+d} \prec_{\mathsf{revlex}} x^{b+d} = x^c$ because $\prec_{\mathsf{revlex}}$ is a term order. Moreover we claim that the vectors in $\mathsf{supp}(d)$ are all sign-compatible with Ua, so $x^{a+d} \notin I_1$: if there were a vector $u \in \mathsf{supp}(d) \subset \mathsf{supp}(c)$ not sign-compatible with $Ua = Ub$, then u would not be sign-compatible with Uc which is a contradiction to $x^c \notin I_1^2$. This concludes the proof that all the monomials in $I_1' \setminus I_1^2$ have the above property.

By 4.3.5 we know that every $\mathbb{Z}$-tension has a representation in $\mathbb{K}[x]/I_1^2$ and the above observation tells us that $\mathbb{K}[x]/(I_1^2 + I_1') = \mathbb{K}[x]/I_1$ contains precisely the $\prec_{\mathsf{revlex}}$-minimal representations of all $\mathbb{Z}$-tensions.

4.3.6. The set of $\mathbb{Z}$-tensions is in bijection with the set of non-zero monomials in $\mathbb{K}[x]/I_1$.

We can now view I_2 as a monomial ideal in $\mathbb{K}[x]/I_1$. Still I_2 contains precisely those non-zero monomials in $\mathbb{K}[x]/I_1$ that correspond to a nowhere zero $\mathbb{Z}$-tension. Thus, we are almost there.

4.3.7. The set of nowhere zero $\mathbb{Z}$-tensions is in bijection with the set of non-zero monomials in $I_2 \subset \mathbb{K}[x]/I_1$.

Square-Free Ideals. All that is left to do is to check that I_1 and I_2 are square free. We have already seen this for I_2 and we have seen this for most generators of I_1, but we have yet to check this for the generators of I_1', i.e. those generators that come from the $\prec_{\mathsf{revlex}}$-construction. This boils down to proving the following statement:

4.3.8. Let $a, b \in \mathbb{N}^{\mathcal{U}}$ be stacking and in particular pairwise sign-compatible such that $Ua = Ub$ and $a \prec_{\mathsf{revlex}} b$. If b contains an entry ≥ 2, then there exist $a', b' \in \mathbb{N}^{\mathcal{U}}$ that are stacking such that $Ua' = Ub'$, $a' \prec_{\mathsf{revlex}} b'$ and $b' < b$.

Proof. Without loss of generality, we can assume $\mathsf{supp}(a) \cap \mathsf{supp}(b) = \emptyset$. For if $v \in \mathsf{supp}(a) \cap \mathsf{supp}(b)$, then $a' := a - e_v$ and $b' := b - e_v$ have the desired properties.

Let $v \in \mathcal{U}$ be the rightmost entry such that $a_v \neq b_v$. As $a \prec_{\mathsf{revlex}} b$ we known that $a_v > b_v$. By the above assumption $b_v = 0$ and all entries of b further to the right are zero as well. So v was used in a stacking representation of $Ua = Ub$ and we have $\mathsf{Max}(Ub) \subset \mathsf{supp}(v) \subset \mathsf{supp}(Ub)$.

Let $u \in \mathcal{U}$ be such that $b_u \geq 2$. Define $b' = b - e_u$ whence $b' < b$ and $\mathsf{supp}(b') = \mathsf{supp}(b)$. As b is stacking, so is b'. Moreover for every $e \in \mathsf{supp}(u)$ we have $|(Ub)_e| \geq 2$. Thus u and $Ub - u = Ub'$ are sign-compatible, $\mathsf{supp}(Ub) = \mathsf{supp}(Ub')$ and by 4.3.3

$$\mathsf{Max}(Ub) = \bigcap_{u \in \mathsf{supp}(b)} \mathsf{supp}(u) = \bigcap_{u \in \mathsf{supp}(b')} \mathsf{supp}(u) = \mathsf{Max}(Ub').$$

Therefore $\mathsf{Max}(Ub') \subset \mathsf{supp}(v) \subset \mathsf{supp}(Ub')$ and by 4.3.5 there is a stacking representation a' of Ub' such that $a'_v \geq 1$. But $b'_v = 0$ and in all entries further to the right b' is zero as well. So in the rightmost entry in which a' and b' differ, a' is bigger than b'. Thus $a' \prec_{\mathsf{revlex}} b'$, which completes the proof. □

Conclusion. The above arguments show the following theorem.

4.3.9. Theorem.

For any graph G, there exists a square-free monomial ideal $I_1 \subset \mathbb{K}[x]$ and a square-free monomial ideal I_2 in $\mathbb{K}[X]/I_1$ such that the Hilbert function $\mathsf{H}_{I_2}(k)$ of I_2 with respect to the standard grading counts the number of $\mathbb{Z}$-tensions t with $||t||_\infty = k$.

By 4.3.1 we obtain as a corollary the following analogue of Steingrímsson's Theorem.

4.3.10. Corollary.

For any graph G, there exists a square-free monomial ideal $I_1 \subset \mathbb{K}[x]$ and a square-free monomial ideal I_2 in $\mathbb{K}[X]/I_1$ such that

$$\mathsf{H}_{I_2}(k) = \theta_G(k+1)$$

for all $k \in \mathbb{N}$, where H_{I_2} denotes the Hilbert function of I_2 with respect to the standard grading and θ_G denotes the integral tension polynomial of G.

This was a long proof. One reason for this is that it is entirely self-contained. We are going to see another much shorter proof of this theorem in Section 4.5, which is shorter, because it makes use of tools developed in the last chapter or quoted from the literature. Developing these tools from scratch, at a similar level of detail, would require at least as many pages.

Moreover, the above argument conveys nicely the flexibility as well as the limitations of the sets of the form $I_2 \setminus I_1$ and it shows how a combinatorial concept such as stacking representations arises naturally from working with these algebraic constraints. Also, it shows the power of term orders to guarantee unique representations. In a way, choosing minimal representations with respect to a total order is the one of the most straightforward approaches one might try, and yet it works simply because the definitions fit together[6].

Arguably the most involved part of the above proof has to do with showing that I_1 and I_2 are square-free. In the proof of 4.3.8, our ability to find a representation a' with $v \in \mathsf{supp}(a')$ was essential, and this necessitated a more thorough treatment of stacking representations in the preceding part of the proof. There also the fact that we are working with reverse lexicographic term orders becomes crucial - we could have done with an arbitrary term order otherwise. We are interested in square-free ideals because they correspond to simplicial complexes, which we will deal with next, in Section 4.4. What these have to do with reverse lexicographic term orders and stacking representations, we will examine in Section 4.6.

In Sections 4.4 and 4.5 we are going to take an entirely different approach to obtaining analogues of Steingrímsson's Theorem, that does not use the material of Sections 4.2 and 4.3 at all. This approach will allow us to capture all five counting functions introduced in Section 3.2 as Hilbert functions of Steingrímsson's type.

4.4. Hilbert equals Ehrhart

In this section we relate Ehrhart functions of certain complexes to Hilbert functions of ideals defined in terms of these complexes. We begin with the well-known relation between simplicial complexes and the corresponding Stanley-Reisner ideals, move on to relative simplicial complexes and relative Stanley-Reisner ideals before we finally consider relative polytopal complexes.[7] As Ehrhart functions are defined in terms of geometric simplicial complexes

[6] Note that we did not have to use any deep theorems about term orders!

[7] We introduce the polytopal Stanley-Reisner ideals corresponding to polytopal complexes only later in Section 4.6.

while Stanley-Reisner ideals are defined in terms of abstract or combinatorial simplicial complexes, all the complexes we consider live in both worlds. A geometric simplicial complex Δ has an abstract simplicial complex $\mathsf{comb}(\Delta)$ associated with it, see Section 1.1. However we will not always use the notation comb to distinguish between the two.

Let Δ be an abstract simplicial complex on the ground set V. We identify the elements of the ground set of Δ with the variables in the polynomial ring $\mathbb{K}[x_v : v \in V] =: \mathbb{K}[x]$. Thus sets $S \subset V$ correspond to square-free monomials in $\mathbb{K}[x]$. The **Stanley-Reisner ideal** I_Δ of Δ is generated by the monomials corresponding to the minimal non-faces of Δ, more precisely

$$I_\Delta := \langle x^u \in K[x] | \mathsf{supp}(u) \notin \Delta \rangle.$$

Then, the **Stanley-Reisner ring** of Δ is the quotient $\mathbb{K}[\Delta] = \mathbb{K}[x]/I_\Delta$. We equip the ring $\mathbb{K}[x]$ with the standard grading, that is for any monomial $x^u \in \mathbb{K}[x]$ we have $\deg(x^u) = ||u||_1 = \sum_{i=1}^n u_i$. The fundamental result about Stanley-Reisner rings is this.

4.4.1. Theorem. [Sta96]

Let Δ be a d-dimensional simplicial complex with f_i faces of dimension i for $0 \leq i \leq d$. Then the Hilbert function $\mathsf{H}_{\mathbb{K}[\Delta]}$ of the Stanley-Reisner ring $\mathbb{K}[\Delta]$ satisfies

$$\mathsf{H}_{\mathbb{K}[\Delta]}(k) \quad = \quad \sum_{i=0}^{d} f_i \binom{k-1}{i} \tag{4.4}$$

for $k > 0$ and $\mathsf{H}_{\mathbb{K}[\Delta]}(0) = 1$.

We remark that (4.4) evaluated at zero gives $\sum_{i=0}^{d} f_i \binom{-1}{i} = \chi(\Delta)$, the Euler characteristic of Δ.

If we are given a geometric simplicial complex Δ we will generally use $\mathsf{vert}(\Delta)$ as the ground set of the abstract simplicial complex $\mathsf{comb}(\Delta)$ and identify the variables of $\mathbb{K}[x]$ with the vertices of Δ. In this case we use I_Δ to refer to $I_{\mathsf{comb}(\Delta)}$ and similarly for $\mathbb{K}[\Delta]$.

Now the Ehrhart functions of a unimodular d-dimensional lattice simplex σ^d and its relative interior $\mathsf{relint}\,\sigma^d$ are, respectively,

$$\mathsf{L}_{\sigma^d}(k) = \binom{k+d}{d} \quad \text{and} \quad \mathsf{L}_{\mathsf{relint}\,\sigma^d}(k) = \binom{k-1}{d}. \tag{4.5}$$

Taken together, (4.4) and (4.5) tell us that for any (geometric) simplicial complex Δ in which all simplices are unimodular, the Ehrhart function $\mathsf{L}_\Delta(k) := \#\mathbb{Z}^d \cap k \bigcup \Delta$ of Δ satisfies

$$\mathsf{L}_\Delta(k) \;=\; \textstyle\sum_{\sigma \in \Delta} \mathsf{L}_{\mathsf{relint}\,\sigma}(k) = \sum_{i=0}^{d} f_i \binom{k-1}{i} \;=\; \mathsf{H}_{\mathbb{K}[\Delta]}(k) \tag{4.6}$$

for all $k > 0$. Simply put: the Ehrhart function of a unimodular geometric simplicial complex and the Hilbert function of the corresponding Stanley-Reisner ring coincide. This fact is well-known, see for example [MS05]. Taking the above approach and calculating the Ehrhart functions of open simplices, however, allows us to do without Möbius inversion.

For our purpose we need a more general concept than that of a Stanley-Reisner ring. For an abstract simplicial complex Δ the Hilbert function $\mathsf{H}_{\mathbb{K}[\Delta]}(k)$ counts all those monomials x^u of degree k with $\mathsf{supp}(u) \in \Delta$. We are interested in a pair of simplicial complexes $\Delta' \subset \Delta$, the former being a subcomplex of the latter, and want to count those monomials x^u such that $\mathsf{supp}(u) \notin \Delta'$ but $\mathsf{supp}(u) \in \Delta$. To that end we follow Stanley [Sta96] in calling a pair of simplicial complexes $\Delta' \subset \Delta$ a **relative simplicial complex**. We denote by $I_{\Delta/\Delta'}$ the ideal in $\mathbb{K}[\Delta]$ generated by all monomials x^u with $\mathsf{supp}(u) \notin \Delta'$. We call this the **relative Stanley-Reisner ideal**. We may view $I_{\Delta/\Delta'}$ as a $\mathbb{K}[\Delta]$-module or as a $\mathbb{K}[x]$ module. Either way, its Hilbert function $\mathsf{H}_{I_{\Delta/\Delta'}}(k)$ counts the number of non-zero monomials x^u of degree k in $I_{\Delta'} \setminus I_{\Delta}$ or, equivalently, the number of non-zero monomials x^u in $\mathbb{K}[x]$ with $\mathsf{supp}(u) \in \Delta \setminus \Delta'$. (Notice how the roles of Δ and Δ' swap, depending on whether we formulate the condition using ideals or using complexes). Now, as Stanley remarks, Theorem 4.4.1 carries over to the relative case.

4.4.2. Theorem. [Sta96]

Let $\Delta' \subset \Delta$ be a relative d-dimensional simplicial complex and let f_i denote the number of i-dimensional simplices in $\Delta \setminus \Delta'$. Then

$$\mathsf{H}_{I_{\Delta/\Delta'}}(k) \quad = \quad \sum_{i=0}^{d} f_i \binom{k-1}{i} \tag{4.7}$$

for $k > 0$.

If Δ is a geometric simplicial complex and $\Delta' \subset \Delta$ a subcomplex, we also call the pair $\Delta' \subset \Delta$ a relative geometric simplicial complex and define its relative Stanley-Reisner ideal $I_{\Delta/\Delta'}$ to be $I_{\mathsf{comb}(\Delta)/\mathsf{comb}(\Delta')}$.

By the same argument as above, we conclude that for any (geometric) relative d-dimensional simplicial complex $\Delta' \subset \Delta$, all faces of which are unimodular,

$$\mathsf{L}_{\bigcup\Delta\setminus\bigcup\Delta'}(k) = \sum_{\sigma\in\Delta\setminus\Delta'} \mathsf{L}_{\mathsf{relint}\,(\sigma)}(k) = \sum_{i=0}^{d} f_i \binom{k-1}{i} = \mathsf{H}_{I_{\Delta/\Delta'}}(k) \tag{4.8}$$

for all $k > 0$, i.e. the Ehrhart function of a relative simplicial complex with unimodular faces and the Hilbert function of the associated relative Stanley-Reisner ideal coincide. Moreover this function is a polynomial in k as

$$\binom{k-1}{i} = \frac{1}{i!}\prod_{i=1}^{i}(k-i)$$

is a polynomial for every $i \in \mathbb{Z}_{\geq 0}$ using the convention that $i! = \prod_{j=1}^{i} j$ and empty products are 1.

To be able to deal with the applications in Section 4.5 we need to go one step further. The complexes we will be dealing with, are not going to be simplicial. Their faces will be polytopes. So we define a **relative polytopal complex** to be a pair $\mathcal{C}' \subset \mathcal{C}$ of polytopal complexes, the former a subcomplex of the latter. Our goal is to realize the Ehrhart function $\mathsf{L}_{\bigcup \mathcal{C} \setminus \bigcup \mathcal{C}'}(k)$ as the Hilbert function of a relative Stanley-Reisner ideal.

By the above arguments, it would suffice to require that $\mathcal{C}$ has a unimodular triangulation. But for the sake of convenience we would like to impose a condition on $\mathcal{C}$ that can be checked one face at a time. Requiring that each face of $\mathcal{C}$ has a unimodular triangulation would not be sufficient. A unimodular triangulation Δ_F for each face $F \in \mathcal{C}$ does not guarantee that $\bigcup_{F \in \mathcal{C}} \Delta_F$ is a unimodular triangulation of $\mathcal{C}$: It may be that for faces $F_1, F_2 \in \mathcal{C}$ that share a common face $F = F_1 \cap F_2$ the unimodular triangulations Δ_{F_1} and Δ_{F_2} do not agree on F, i.e.

$$\{F \cap f_1 \mid f_1 \in \Delta_{F_1}\} \neq \{F \cap f_2 \mid f_2 \in \Delta_{F_2}\}.$$

Fortunately there is the notion of a compressed polytope: it suffices to require of each face $F \in \mathcal{C}$ individually that F is compressed, to guarantee that $\mathcal{C}$ as a whole has a unimodular triangulation.

In the following the polytopal complex **generated by** a collection of polytopes S, is the set of all faces F of the polytopes $P \in S$. In other words the polytopal complex generated by S is obtained by closing S under the operation of passing to faces. Recall that a lattice polytope is empty if it contains no lattice points except its vertices.

Let $P \subset \mathbb{R}^d$ be a lattice polytope. Let $\prec$ be a total ordering of the lattice points in P. The **pulling triangulation** $\mathsf{pull}(P; \prec)$ of P with respect to the total ordering $\prec$ is defined recursively as follows. If P is an empty simplex, then $\mathsf{pull}(P; \prec)$ is the complex generated by P. Otherwise $\mathsf{pull}(P; \prec)$ is the complex generated by the set of polytopes

$$\bigcup_F \{\mathsf{conv}\{v, G\} : G \in \mathsf{pull}(F; \prec)\}$$

where v is the $\prec$-minimal lattice point in P and the union runs over all faces F of P that do not contain v. See also Sturmfels [Stu96]. This construction yields a triangulation and the vertices of $\mathsf{pull}(P, \prec)$ are lattice points in P. Pulling triangulations need not be unimodular, in fact the simplices in a pulling triangulation do not even have to be empty! Polytopes whose pulling triangulation is always unimodular get a special name. A polytope P is **compressed** if for any total ordering $\prec$ on the vertex set the pulling triangulation $\mathsf{pull}(P, \prec)$ is unimodular. These definitions have the following well-known properties.

4.4.3.

1. Any $\dim(P)$-dimensional simplex in $\mathsf{pull}(P, \prec)$ contains the $\prec$-minimal lattice point v in P as a vertex.
2. $\mathsf{pull}(F, \prec) = \mathsf{pull}(P, \prec) \cap F$ for any total order $\prec$ on the lattice point in P and any face F of P.
3. All faces of a compressed polytopes are compressed.
4. Compressed polytopes are empty.

Proof. 1. and 2. follow directly from the definition. As all faces of a unimodular simplex are unimodular 2. implies 3. To show 4. we argue as follows. Suppose that, using our definition, a lattice polytope P contains a lattice point $z \in \mathsf{relint}\, P$ in its relative interior. If we now define $\prec$ such that a vertex of P is $\prec$-minimal, then there will be some simplex $\sigma \in \mathsf{pull}(P, \prec)$ that contains z, but z is not a vertex of σ. Thus σ is not unimodular and P is not compressed. So no compressed lattice polytope contains a lattice point in its relative interior. Combining this with 3. completes the proof. □

In some sources [OH01], [Sul04] the pulling triangulation of a polytope is defined not with respect to all lattice points in P, but only with respect to the vertices of P. It is important to note that this gives rise to the same concept of a compressed polytope. Let us call the polytopes whose pulling triangulations with respect to this alternative definition are unimodular "vertex compressed". Suppose P is vertex compressed. Then in particular it is empty[8]. So the two notions of pulling triangulation coincide. On the other hand if P is compressed, then by 4.4.3 we also know that P is empty and again the two notions of pulling triangulation coincide. So the notions of "compressed" and "vertex compressed" are equivalent.

For more information on pulling triangulations and compressed polytopes we refer to [Stu96], [OH01], [Sul04], [Stu91] and [LRS09].

Now, if $\mathcal{C}$ is a polytopal complex with integral vertices such that every face $P \in \mathcal{C}$ is compressed, then we can fix an arbitrary total order $\prec$ on $\bigcup \mathcal{C} \cap \mathbb{Z}^d$ and construct the pulling triangulations $\mathsf{pull}(P, \prec)$ of all faces $P \in \mathcal{C}$ with respect to that one global order $\prec$. By 4.4.3 this means that for any two $P_1, P_2 \in \mathcal{C}$ that share a common face $F = P_1 \cap P_2$ the triangulations induced on F agree: $\mathsf{pull}(P_1, \prec) \cap F = \mathsf{pull}(F, \prec) = \mathsf{pull}(P_2, \prec) \cap F$. Thus $\mathsf{pull}(\mathcal{C}, \prec) := \bigcup_{F \in \mathcal{C}} \mathsf{pull}(F, \prec)$ is a unimodular triangulation of $\mathcal{C}$ with $\mathsf{vert}(\mathcal{C}) = \mathsf{vert}(\mathsf{pull}(\mathcal{C}, \prec))$. We abbreviate $\Delta := \mathsf{pull}(\mathcal{C}, \prec)$.

If $\mathcal{C}'$ is any subcomplex of $\mathcal{C}$, we define Δ' to be the subcomplex of Δ consisting of those faces $F \in \Delta$ such that $F \subset \bigcup \mathcal{C}'$. So

$$\mathsf{L}_{\bigcup \mathcal{C} \setminus \bigcup \mathcal{C}'}(k) = \mathsf{L}_{\bigcup \Delta \setminus \bigcup \Delta'}(k) = \mathsf{H}_{I_{\Delta/\Delta'}}(k) \tag{4.9}$$

for $k > 0$ which means that we have realized the Ehrhart function of $\bigcup \mathcal{C} \setminus \bigcup \mathcal{C}'$ as the Hilbert function of the relative Stanley-Reisner ideal $I_{\Delta/\Delta'}$. Moreover, we have already seen that this function is a polynomial. We summarize these results in the following theorem.

[8]By definition only the vertices of P can be vertices of $\mathsf{pull}(P, \prec)$.

4.4.4. Theorem.

Let $\mathcal{C}$ be a polytopal complex. If all faces of $\mathcal{C}$ are compressed lattice polytopes, then for any subcomplex $\mathcal{C}' \subset \mathcal{C}$ there exists a relative Stanley-Reisner ideal $I_{\Delta,\Delta'}$ such that for all $0 < k \in \mathbb{N}$

$$\mathsf{L}_{\bigcup \mathcal{C} \setminus \bigcup \mathcal{C}'}(k) = \mathsf{H}_{I_{\Delta/\Delta'}}(k)$$

and this function is a polynomial.

In particular, the Ehrhart function of any inside-out polytope whose full-dimensional cells are compressed lattice polytopes is a Hilbert function of a relative Stanley-Reisner ideal.

4.5. Counting Polynomials as Hilbert Functions (Geometrically)

We now want to apply Theorem 4.4.4 to obtain analogues of Steingrímsson's Theorem 4.2.1 for all five of our counting polynomials from Section 3.2. We know realizations of all of these as Ehrhart functions, so all we have to do to apply Theorem 4.4.4 is to show that their faces are compressed lattice polytopes. Our tool for proving this is going to be the following theorem by Ohsugi and Hibi which states that lattice polytopes that are slices of the unit cube are automatically compressed.[9]

4.5.1. Theorem. (Ohsugi and Hibi [OH01, Theorem 1.1])

Let P be a lattice polytope in $\mathbb{R}^n$. If P is lattice isomorphic to the intersection of an affine subspace with the unit cube, i.e. $P \approx [0,1]^n \cap L$ for some affine subspace L, then P is compressed.

Sullivant [Sul04] showed that this property even characterizes compressed lattice polytopes! We are not going to use that fact here, however. The Theorem of Ohsugi and Hibi is all we need.

Integral Tension Polynomial

Let $G = (V, E)$ be a graph and T a spanning multi-tree. Let C be the corresponding cycle matrix of G. As we have seen in Section 3.4, ker C is the tension space of G and the lattice points in ker C are in bijection with the integral tensions of G. The k-tensions of G are those lattice points t in the tension space with $||t||_\infty < k$, in other words, the k-tensions are the lattice points in ker $C \cap (-k, k)^E$. The nowhere zero k-tensions, then, are those lattice

[9] Actually, Ohsugi and Hibi showed a more general result, but this will suffice for our purposes.

points $t \in \ker C \cap (-k,k)^E$ with $t(e) \neq 0$ for all $e \in E$. So consider the open polytope $\ker C \cap (-1,1)^E$ and let

$$\mathsf{Coord} = \Big\{ \{t \in \mathbb{R}^E : t(e) = 0\} \ : \ e \in E \Big\}$$

denote the hyperplane arrangement consisting of all coordinate hyperplanes. Then $(\ker C \cap (-1,1)^E, \mathsf{Coord})$ is an (open) inside-out polytope, and we have just sketched a proof that the Ehrhart function $\mathsf{L}_{\ker C \cap (-1,1)^E \setminus \bigcup \mathsf{Coord}}$ of this inside-out polytope is the integral tension polynomial. Details of this proof can be found in [Dal08], where this result is used to prove the Integral Tension Reciprocity Theorem 3.3.4. We summarize:

4.5.2. $\mathsf{L}_{\ker C \cap (-1,1)^E \setminus \bigcup \mathsf{Coord}}(k) = \theta_G(k)$.

Note that we are in fact dealing with the integral tension polynomial $\theta_G(k)$ and not with the shifted polynomial $\theta_G(k+1)$! This is due to the fact that we are dealing with the open polytope $\ker C \cap (-1,1)^E$ which does not include the $0, \pm 1$-tensions except the all zero tension.

Of course, this inside-out polytope can also be viewed as a relative polytopal complex $\mathcal{C}' \subset \mathcal{C}$. $\mathcal{C}$ consists of the closures of the cells of $(\ker C \cap (-1,1)^E, \mathsf{Coord})$ and all their respective faces and $\mathcal{C}'$ is the collection of all faces of $\mathcal{C}$ that lie on the boundary of $\ker C \cap [-1,1]^E$ or on one of the coordinate hyperplanes. More precisely, we let $\mathsf{comp}(\mathsf{Coord})$ denote the complex given by the coordinate hyperplane arrangement and define

$$\begin{aligned} \mathcal{C} &= \{F \cap \ker C \cap [-1,1]^E \ : \ F \in \mathsf{comp}(\mathsf{Coord})\}, \\ \mathcal{C}' &= \{F \in \mathcal{C} \ : \ F \subset \partial[-1,1]^E \cup \bigcup \mathsf{Coord}\}. \end{aligned}$$

So we have captured the integral tension polynomial as the Ehrhart function of a relative polytopal complex.

4.5.3. $\mathsf{L}_{\bigcup \mathcal{C} \setminus \bigcup \mathcal{C}'}(k) = \theta_G(k)$.

We now wish to apply Theorem 4.4.4. To this end we have to show that all the facets of $\mathcal{C}$ are compressed lattice polytopes. We know by Theorem 3.5.4 that C is totally unimodular. Because every facet $F \in \mathcal{C}$ is of the form $\ker C \cap \prod_{e \in E}[l_e, l_e+1]$ where $l_e \in \{0, -1\}$ for all $e \in E$, we know by the Theorem of Hoffman and Kruskal 3.5.1 that F is a lattice polytope. Because F is a lattice polytope of the form $\ker C \cap \prod_{e \in E}[l_e, l_e+1]$ we know by the Theorem of Ohsugi and Hibi 4.5.1 that F is compressed. Thus we can apply Theorem 4.4.4 and obtain

4.5.4. Theorem.

For any graph G there exists a relative Stanley-Reisner ideal $I_{\Delta,\Delta'}$ such that for all $0 < k \in \mathbb{N}$

$$\theta_G(k) = \mathsf{H}_{I_{\Delta/\Delta'}}(k).$$

This improves on our earlier result Corollary 4.3.10 insofar as we now obtain the integral tension polynomial $\theta_G(k)$ itself and not the shifted polynomial $\theta_G(k+1)$. We could obtain $\theta_G(k+1)$ by excluding the faces of $\mathcal{C}$ that lie on $\partial[-1,1]^E$ but not on $\bigcup \mathsf{Coord}$ from $\mathcal{C}'$. Note that Theorem 4.3.9 is a stronger result insofar as there we constructed a Hilbert function that counted the nowhere zero $k+1$-tensions with maximal edge weight *exactly* $k+1$. We can obtain Theorem 4.3.9 with our geometric method using the relative complex $\mathcal{C}_1' \subset \mathcal{C}_1$ where $\mathcal{C}_1$ consist of those faces of $\mathcal{C}$ that lie on the boundary $\partial[-1,1]^E$ of the cube and where $\mathcal{C}_1'$ consist of those faces of $\mathcal{C}_1$ that lie on a coordinate hyperplane.

Integral Flow Polynomial

Let $G = (V, E)$ be a graph. To capture the integral flow polynomial of G as a Hilbert function we proceed exactly as in the case of the integral tension polynomial. The inside-out polytope whose Ehrhart function is the integral tension polynomial is $(\ker A \cap (-1,1)^E, \mathsf{Coord})$ where A denotes the incidence matrix of G. This is precisely the inside-out polytope studied in [BZ06b]. Our relative polytopal complex $\mathcal{C}' \subset \mathcal{C}$ is the one determined by the inside-out polytope: $\mathcal{C}$ consists of the closures of the cells of $(\ker A \cap (-1,1)^E, \mathsf{Coord})$ and their respective faces while $\mathcal{C}'$ contains all those faces of $\mathcal{C}$ that are contained in $\bigcup \mathsf{Coord}$. Given these definitions

4.5.5. $\mathsf{L}_{\bigcup\mathcal{C}\setminus\bigcup\mathcal{C}'}(k) = \varphi_G(k).$

Again we need to check that the facets of $\mathcal{C}'$, that is the closures of the cells of $(\ker A \cap (-1,1)^E, \mathsf{Coord})$, are compressed lattice polytopes. The same method as in the case of the integral tensions works. A facet F of $\mathcal{C}$ has the form $\ker A \cap \prod_{e\in E}[l_e, l_e+1]$ where $l_e \in \{0, -1\}$ for all $e \in E$. Thus, by the Theorem of Hoffman and Kruskal 3.5.1 F is a lattice polytope and by the Theorem of Ohsugi and Hibi 4.5.1 is compressed. Therefore, by Theorem 4.4.4:

4.5.6. Theorem.

For any graph G there exists a relative Stanley-Reisner ideal $I_{\Delta,\Delta'}$ such that for all $0 < k \in \mathbb{N}$

$$\varphi_G(k) = \mathsf{H}_{I_{\Delta/\Delta'}}(k).$$

Modular Tension Polynomial

In Section 3.9 we have already met two relative complexes whose Ehrhart function is the modular tension polynomial of a graph $G = (V, E)$. Here we make use of the first construction. Let C be the cycle matrix of G with respect to any spanning multi-tree T. We define $P_d^\circ = \{t \in \mathbb{Z}^E | 0 < t < 1, Ct = d\}$ and $P_d = \{t \in \mathbb{Z}^E | 0 \leq t \leq 1, Ct = d\}$ for $d \in \mathcal{D}_G := \{d \in \mathbb{Z}^{|E|-|V|+c} : P_d^\circ \neq \emptyset\}$ and recall 3.9.1 which stated

$$\bar{\theta}_G(k) = \sum_{d \in \mathcal{D}_G} \mathsf{L}_{P_d^\circ}(k).$$

Now we define $\mathcal{C}$ to be the complex given by the P_d for $d \in \mathcal{D}_G$ and their respective faces and we let $\mathcal{C}'$ denote the subcomplex consisting of all those faces $F \in \mathcal{C}$ that lie in the boundary of some P_d, that is $F \subset \partial P_d$ for some $d \in \mathcal{D}_G$. Then

4.5.7. $\bar{\theta}_G(k) = \mathsf{L}_{\bigcup \mathcal{C} \setminus \bigcup \mathcal{C}'}(k)$.

As we have already seen in Section 3.9 the facets P_d of $\mathcal{C}$ are lattice polytopes by the Theorem of Hoffman and Kruskal 3.5.1. Moreover they have the form

$$P_d = \{t \mid Ct = d\} \cap [0,1]^E$$

so by the Theorem of Ohsugi and Hibi 4.5.1 all P_d are compressed. Thus, by Theorem 4.4.4:

4.5.8. Theorem.

For any graph G there exists a relative Stanley-Reisner ideal $I_{\Delta,\Delta'}$ such that for all $0 < k \in \mathbb{N}$

$$\bar{\theta}_G(k) = \mathsf{H}_{I_{\Delta/\Delta'}}(k).$$

Modular Flow Polynomial

As with the modular tension we have already seen two relative complexes (in Sections 3.6 and 3.7, respectively) whose Ehrhart functions are the modular flow polynomial of a given graph $G = (V, E)$. Again we make use of the first construction. Let A denote the incidence matrix of G. We define $P_b^\circ := \{f \in \mathbb{R}^E \mid 0 < f < 1, Af = b\}$ and $P_b := \{f \in \mathbb{R}^E \mid 0 \leq f \leq 1, Af = b\}$ for $b \in \mathcal{B}_G := \{b \in \mathbb{Z}^V | P_b^\circ \neq \emptyset\}$ and recall 3.6.2 which stated

$$\bar{\varphi}(k) = \sum_{b \in \mathcal{B}_G} \mathsf{L}_{P_b^\circ}(k).$$

We define the polytopal complex $\mathcal{C}$ to consist of the polytopes P_b for $b \in \mathcal{B}_G$ and their respective faces. The subcomplex $\mathcal{C}'$ consists of those faces $F \in \mathcal{C}$ that are contained in the boundary of some P_b, that is $F \subset \partial P_b$ for some $b \in \mathcal{B}_G$.

As we have already seen in Section 3.6 the facets P_b of $\mathcal{C}$ are lattice polytopes by the Theorem of Hoffman and Kruskal 3.5.1. Moreover they have the form

$$P_d = \{f \mid Af = b\} \cap [0,1]^E$$

so by the Theorem of Ohsugi and Hibi 4.5.1 any P_b is compressed. Thus, by Theorem 4.4.4:

4.5.9. Theorem.

For any graph G there exists a relative Stanley-Reisner ideal $I_{\Delta,\Delta'}$ such that for all $0 < k \in \mathbb{N}$

$$\bar{\varphi}_G(k) = \mathsf{H}_{I_{\Delta/\Delta'}}(k).$$

Chromatic Polynomial

Finally we turn to the chromatic polynomial to recapture and improve upon Steingrímsson's Theorem 4.2.1. Let $G = (V, E)$ be a graph. First of all we exclude the case of graphs G that have a loop. Those graphs do not have any proper colorings, so their chromatic polynomial is constant zero which can be realized as the Hilbert function of the ring $\mathbb{K}[x]/\mathbb{K}[x]$. Note that $\mathbb{K}[x] = \langle 1 \rangle$ is a square-free monomial ideal. So we assume that G does not have loops.

A coloring $c : V \to \mathbb{Z}$ is an element of the vertex space $\mathbb{Z}^V$. c is proper if $c(u) \neq c(v)$ for any edge $e = uv \in E$, that is if the lattice point c does not lie on the hyperplane $H_{uv} = \{c \in \mathbb{R}^V \mid c(u) - c(v) = 0\}$. So let

$$\mathsf{Braid} = \{H_{uv} \;:\; u, v \in V, u \neq v\}$$

denote the arrangement consisting of all theses hyperplanes. This arrangement is called the **braid arrangement** or the Coxeter arrangement of type $A_{|V|}$. Let $\mathcal{H} \subset \mathsf{Braid}$ denote the subarrangement consisting of all those hyperplanes H_{uv} with $e = uv \in E$. As G does not have loops, there are no degenerate hyperplanes in $\mathcal{H}$. The lattice points in $\mathbb{R}^V \setminus \bigcup \mathcal{H}$ correspond to the proper colorings of G.

How do we model k-colorings? Of the many possible choices, we take the following: we take our k colors to be the integers $0, \ldots, k-1$. Thus the k-colorings of G are in bijection with the lattice points in the half-open cube $[0,k)^V$ and by the above arguments the proper k-colorings of G are in bijection with the lattice points in $[0,k)^V \setminus \bigcup \mathcal{H}$. This modeling is slightly different from the one given in [BZ06a], who identify k-colorings with lattice points in $(0, k+1)^V \setminus \bigcup \mathcal{H}$. We have now identified the chromatic polynomial of G with the Ehrhart function of the "half-open" inside-out polytope $([0,k)^V, \mathcal{H})$.

We proceed to define our relative polytopal complex $\mathcal{C}' \subset \mathcal{C}$ as follows. We let $\mathcal{C}$ be the complex given by the closures of the connected components of $[0,1)^V \setminus \bigcup \mathsf{Braid}$ and their respective faces. We let $\mathcal{C}'$ consist of those faces $F \in \mathcal{C}$ that have at least one of the following two properties: 1) F is contained in a hyperplane $H_{uv} \in \mathcal{H}$. 2) There exists a $v \in V$ such that $c(v) = 1$ for all $c \in F$. This means that F is contained in one of the "open" faces of the half-open cube $[0,1)^V$. This relative complex $\mathcal{C}' \subset \mathcal{C}$ now has the property that $k(\bigcup \mathcal{C} \setminus \bigcup \mathcal{C}') = [0,k)^V \setminus \bigcup \mathcal{H}$ and thus

4.5.10. $\mathsf{L}_{\bigcup \mathcal{C} \setminus \bigcup \mathcal{C}'}(k) = \chi_G(k)$.

Again we have to check that the facets of $\mathcal{C}$ are compressed lattice polytopes. It is a well-known fact that the facets of $\mathcal{C}$ are unimodular lattice simplices. This fact can be seen as follows. Let p, q be two vertices of $[0,1]^V$. There is a facet containing both p and q if and only if p and q are not separated by one of the hyperplanes H_{uv}. p and q are separated by H_{uv} if and only if $p(u) - p(v)$ and $q(u) - q(v)$ have opposite sign, say $p(u) - p(v) < 0$ and $q(u) - q(v) > 0$. As p, q are $0,1$-vectors this is equivalent to $p(v) = q(u) = 1$ and $p(u) = q(v) = 0$, which just means that p and q are incomparable in the dominance order $\leq$. Thus we have seen that two vertices p and q of the cube lie on a common face of $\mathcal{C}$ if and only if p and q are comparable in the dominance order. The vertex sets of the facets of $\mathcal{C}$ are thus maximal chains in $\{0,1\}^V$ under the dominance order, i.e. the maximal chains in the Boolean lattice. These chains are of cardinality $|V| + 1$ and the vertices in such a chain are affinely independent, so the facets of $\mathcal{C}$ are lattice simplices. Let M be the matrix having the vertices in such a chain as columns, with the exception of the vertex 0. Given a suitable permutation of rows and columns, this matrix is triangular with ones on the diagonal. Thus $\det M = 1$ and the simplex is unimodular.

So the facets are unimodular lattice simplices. As such these are clearly compressed, for every pulling triangulation is just the simplex itself - and this simplex is unimodular. Therefore we can apply Theorem 4.4.4.

4.5.11. Theorem.

For any graph G there exists a relative Stanley-Reisner ideal $I_{\Delta,\Delta'}$ such that for all $0 < k \in \mathbb{N}$

$$\chi_G(k) = \mathsf{H}_{I_{\Delta/\Delta'}}(k).$$

This is an improvement upon Steingrímsson's Theorem insofar as we obtain the chromatic polynomial $\chi_G(k)$ itself as a Hilbert function of a relative Stanley-Reisner ideal and not the shifted polynomial $\chi_G(k+1)$. To obtain the shifted chromatic polynomial using the above construction we would need to consider the closed cube $[0,1]^V$ instead of the half-open cube $[0,1)^V$.

4.6. *The Combinatorial versus the Geometric Approach*

In this section we reexamine the combinatorial proof we gave in Section 4.3 of the fact that the tension polynomial is a Hilbert function. We relate it to the geometric methods from Sections 4.4 and 4.5 and show that it can be generalized to yield a combinatorial/algebraic proof of Theorem 4.4.4.

Let us go through the combinatorial proof of 4.3.9 and reconsider each step geometric terms. Keep in mind, though, that in 4.3.9 our goal was to establish a *bijection* between the monomials in I_2 and the nowhere zero $\mathbb{Z}$-tensions that restricts to a bijection between the monomials in I_2 of degree k and the nowhere zero $\mathbb{Z}$-tensions with maximal weight *equal to* k. In 4.5.4 our goal was to give a *surjection* from the monomials in $I_{\Delta/\Delta'}$ onto the nowhere zero $\mathbb{Z}$-tensions that restricts to a bijection between the monomials in $I_{\Delta/\Delta'}$ of degree k onto the nowhere zero $\mathbb{Z}$-tensions with maximal absolute weight *less than* k.

As our first step in Section 4.3 we established a bijection between variables in $\mathbb{K}[x]$ and lattice points on the boundary of $\ker C \cap [-1,1]^E$. Multiplication of variables corresponded to addition of vectors.

Next we defined the ideal

$$I_1^1 := \langle x_u x_v \mid u \text{ and } v \text{ are not sign-compatible} \rangle.$$

The purpose was to make sure that whenever we have $x^a, x^b \in \mathbb{K}[x]/I_1^1$ such that Ua and Ub are nowhere zero, then $a+b$ is also nowhere zero. We can ensure this by requiring Ua and Ub to lie in the same orthand which is just what requiring u and v to be sign compatible achieves. Note that the orthands of $\mathbb{R}^E$ form a fan, which we call the **orthand fan**. I_1^1 now encodes the constraint that we may only add vectors that lie in the same cone of the orthand fan.

Moreover we wanted that $\deg(x^a) = ||Ua||_\infty$ for any $x^a \in \mathbb{K}[x]/I_1$. To guarantee this we needed to require in addition that all $u \in \mathsf{supp}(a)$ share a common element $e \in \mathsf{Max}(u)$. This just means that all $u \in \mathsf{supp}(a)$ lie on one face of the cube $[-1,1]^E$. Thus we called a vector a a stacking representation if the $u \in \mathsf{supp}(a)$ are pairwise sign-compatible and $\bigcap_{u \in \mathsf{supp}(a)} \mathsf{Max}(u) \neq \emptyset$ and defined

$$I_1^2 := \langle x^a \mid a \text{ is not a stacking representation} \rangle.$$

Note that for any set of vectors $u_1, \ldots, u_l$ the conditions

$$\bigcap_{i=1}^{l} \mathsf{Max}(u) \neq \emptyset \quad \text{and} \quad \mathsf{sgn}(u_i(e)) = \mathsf{sgn}(u_j(e)) \text{ for any } 1 \le i,j \le l \text{ and } e \in \bigcap_{i=1}^{l} \mathsf{Max}(u)$$

mean that the vectors u_i all lie in the same cone of the face fan of the cube $[-1,1]^E$. Which is cone of the face fan that is, is given by $\bigcap_{i=1}^{l} \mathsf{Max}(u)$. If $\mathcal{F}$ denotes the intersection of the orthand fan and the face fan of the cube, then a stacking representation a is a representation such that all $u \in \mathsf{supp}(a)$ lie in a common cone of $\mathcal{F}$ and I_1^2 encodes the constraint that all representations are stacking.

So, finding a stacking representation of an arbitrary $\mathbb{Z}$-tension t just means finding a non-negative integral representation of t in terms of $0, \pm 1$-tensions that lie in the inclusion minimal cone of $\mathcal{F}$ that contains t. In terms of complexes this can be phrased as follows. Let $\mathcal{C}$ be the intersection of $\partial[-1,1]^E \cap \ker C$ and $\mathcal{F}$. Given any $\mathbb{Z}$-tension t, we find the inclusion-minimal face $F \in \mathcal{C}$ such that $t \in kF$. Such faces exists! Then, a stacking representation of t is an integral non-negative representation of t in terms of lattice points in F. Now 4.3.5 simply stated that for any face $F \in \mathcal{C}$, any $\mathbb{Z}$-tension $t \in kF$ has an integral non-negative representation in terms of lattice points in F *and we can prescribe an arbitrary lattice point in F that has to appear in the representation.*

This motivates the following definition. For any polytopal complex $\mathcal{C}$ in $\mathbb{R}^d$ and any $x \in \mathbb{R}^d$ let $\mathsf{minface}_{\mathcal{C}}(x)$ denote the inclusion minimal face of $\mathcal{C}$ that contains x, if such a face exists. We say $\mathcal{C}$ has **greedy representations** if the following holds:

> For any $k \in \mathbb{N}$ and for any $z \in \mathbb{Z}^d \cap k\mathcal{C}$ the face $\mathsf{minface}_{\mathcal{C}}(\frac{1}{k}z)$ contains a lattice point and for any such lattice point u there exists a non-negative integral representation a of z in terms of lattice points in $\mathsf{minface}_{\mathcal{C}}(\frac{1}{k}z)$ with $||a||_1 = k$ such that $u \in \mathsf{supp}(a)$.[10]

We say a polytope P has greedy representations if P, viewed as a polytopal complex, has. In light of this new concept 4.3.5 states simply that in the case of $\mathbb{Z}$-tensions the complex $\mathcal{C}$ has greedy representations. This more abstract terminology allows us to give a proof of Theorem 4.4.4 in combinatorial terms, simply by transferring what we did in the tension case, as we will see next.

This is not surprising: As it turns out, for a polytope P having greedy representations is equivalent to being a compressed lattice polytope. We will see that at the end of this section. Right away we only make two observations. The first thing to note about this definition is that the first part, namely the condition that "$\mathsf{minface}_{\mathcal{C}}(\frac{1}{k}z)$ contains a lattice point", is equivalent to saying that the vertices of $\mathcal{C}$ are lattice points. The other thing to notice is that the second part implies that a polytope P with greedy representations is empty. For let $z \in \mathbb{Z}^d \cap P$ and $v \in \mathsf{vert}(\mathsf{minface}_P(z))$. Then there exists a non-negative integral representation a of z with $||a||_1 = 1$. But this just means that $z = a \in \mathsf{vert}(P)$.

Next, we need to transfer our construction of the ideals I_1 and I_2 in the tension case to an arbitrary complex $\mathcal{C}$. For any polytopal complex $\mathcal{C}$, we can generalize the notion of a Stanley-Reisner ideal, which was only defined for simplicial complexes, as follows. The **polytopal Stanley-Reisner ideal** $I_{\mathcal{C}}$ is defined by

$$I_{\mathcal{C}} := \langle x^a | \mathsf{supp}(a) \not\subset \mathsf{vert}(F) \text{ for every face } F \in \mathcal{C} \rangle.$$

[10] I.e. there exist $\lambda_v \in \mathbb{Z}_{\geq 0}$ for $v \in \mathbb{Z}^d \cap \mathsf{minface}_{\mathcal{C}}(\frac{1}{k}z)$ such that $\sum_{v \in \mathbb{Z}^d \mathsf{minface}_{\mathcal{C}}(\frac{1}{k}z))} \lambda_v v = z$, $\sum_{v \in \mathbb{Z}^d \mathsf{minface}_{\mathcal{C}}(\frac{1}{k}z))} \lambda_v = k$ and $\lambda_u \geq 1$.

4.6.1. If a polytopal complex $\mathcal{C}$ has greedy representations, then the monomials of degree k not in $I_{\mathcal{C}}$ are mapped surjectively onto the lattice points in $k \bigcup \mathcal{C}$.

Proof. Let x^a be a monomial in $\mathbb{K}[x]$ with $\deg(x^a) = k$ and $x^a \notin I_{\mathcal{C}}$. As Ua is a non-negative combination of the vertices of some face $F \in \mathcal{C}$ and as all of these vertices are integral, $Ua \in kF \cap \mathbb{Z}^d$. Conversely, let $z \in k \bigcup \mathcal{C} \cap \mathbb{Z}^d$. There is a non-negative integral representation a of z in terms of $\mathsf{vert}(\mathsf{minface}_{\mathcal{C}}(\frac{1}{k}z))$. So x^a is a monomial not in $I_{\mathcal{C}}$ such that $Ua = z$. □

However in this argument we have not yet used our ability to force a particular vertex $u \in \mathsf{vert}(\mathsf{minface}_{\mathcal{C}}(\frac{1}{k}z))$ to appear in the representation. This will be of importance later.

Now we are going to use the same construction using the reverse lexicographic term order that we applied in Section 4.3 to turn the above surjection into a bijection. There is one fine point though:

> We require, without loss of generality, that for every $z \in \mathbb{Z}^d$ there is *at most one* $k \in \mathbb{N}$ such that $z \in k \bigcup \mathcal{C}$.

That we can indeed do so without loss of generality is shown by the following argument. Given any polytopal complex $\mathcal{C} \subset \mathbb{R}^d$ we can embed $\mathcal{C}$ in $\mathbb{R}^{d+1}$ at height 1 by passing to $\mathcal{C}' := \{P \times \{1\} \mid P \in \mathcal{C}\}$. That way the k-th dilate of $\mathcal{C}'$ lies in $\mathbb{R}^d \times \{k\}$ and so every point in $\mathbb{Z}^{d+1}$ lies in $k \bigcup \mathcal{C}'$ for at most one k as desired. The important thing is that this change did not affect the Ehrhart function: $\mathsf{L}_{\mathcal{C}}(k) = \mathsf{L}_{\mathcal{C}'}(k)$ for all k. So if we can show $\mathsf{L}_{\mathcal{C}'}(k)$ to be a Hilbert function the same holds for $\mathsf{L}_{\mathcal{C}}(k)$. From now on we assume that $\mathcal{C}$ has the above property.

For a fixed ordering on the vertices of $\mathcal{C}$, we consider the reverse lexicographic term ordering $\prec_{\mathsf{revlex}}$ on $\mathbb{K}[x]$. If I_U denotes the toric ideal

$$I_U := \langle x^a - x^b \mid x^a, x^b \notin I_{\mathcal{C}},\ Ua = Ub \rangle$$

then its initial ideal $\mathsf{in}_{\prec_{\mathsf{revlex}}}(I_U)$ is

$$\mathsf{in}_{\prec_{\mathsf{revlex}}}(I_U) = \langle x^b \mid x^b \notin I_{\mathcal{C}}, \exists x^a \notin I_{\mathcal{C}} : Ua = Ub \wedge x^a \prec_{\mathsf{revlex}} x^b \rangle.$$

Because $\prec_{\mathsf{revlex}}$ is a term order, the monomials in $\mathsf{in}_{\prec_{\mathsf{revlex}}}(I_U)$ are *precisely* those monomials $x^b \notin I_{\mathcal{C}}$ such that there exists an $x^a \notin I_{\mathcal{C}}$ with $Ua = Ub$ and $x^a \prec_{\mathsf{revlex}} x^b$. So every monomial $x^a \notin I_{\mathcal{C}} + \mathsf{in}_{\prec_{\mathsf{revlex}}}(I_U)$ is the $\prec_{\mathsf{revlex}}$-minimal representation of the lattice point Ua. By our assumption that every z is in $k \bigcup \mathcal{C}$ for at most one k, we know that if $Ua \in k \bigcup \mathcal{C}$ then the $\prec_{\mathsf{revlex}}$-minimal representation x^a has $||a||_1 = \deg(x^a) = k$. So not only does every lattice point $\bigcup_{k \in \mathbb{N}} k \bigcup \mathcal{C}$ have a unique representation in $\mathbb{K}[x]/(I_{\mathcal{C}} + \mathsf{in}_{\prec_{\mathsf{revlex}}}(I_U))$ but that representation also has the right degree. We summarize:

4.6.2. If a polytopal complex $\mathcal{C}$ has greedy representations, then the monomials of degree k not in $I_\mathcal{C} + \mathsf{in}_{\prec_{\mathsf{revlex}}}(I_U)$ are mapped bijectively onto the lattice points in $k\bigcup\mathcal{C}$.

We move on to the question, whether $\mathsf{in}_{\prec_{\mathsf{revlex}}}(I_U)$ is square-free. All we have to do is transfer our proof of 4.3.8 into this more abstract setting. Again, we have to show that if $x^a, x^b \notin I_\mathcal{C}$ such that $Ua = Ub$, $x^a \prec_{\mathsf{revlex}} x^b$ and $b_u \geq 2$, then there are $x^{a'}, x^{b'} \notin I_\mathcal{C}$ such that $Ua' = Ub'$, $x^{a'} \prec_{\mathsf{revlex}} x^{b'}$ and $b' < b$. Again we can assume without loss of generality that $\mathsf{supp}(a) \cap \mathsf{supp}(b) = \emptyset$. Let v be the rightmost index such that $a_v > b_v = 0$. Consider $b' := b - e_u < b$. Then $k := ||a||_1 = ||b||_1 = ||b'||_1 + 1$. Using $u \in \mathsf{supp}(a)$, $Ua = Ub$ and $\mathsf{supp}(b') = \mathsf{supp}(b)$ we know that

$$u \in \mathsf{vert}(\mathsf{minface}_\mathcal{C}(\frac{1}{k}Ua)) = \mathsf{vert}(\mathsf{minface}_\mathcal{C}(\frac{1}{k}Ub)) = \mathsf{vert}(\mathsf{minface}_\mathcal{C}(\frac{1}{k-1}Ub')).$$

As $\mathcal{C}$ has greedy representations, we know that there exists a non-negative integral representation a' of Ub' with $u \in \mathsf{supp}(a')$. But this means that $x^{a'} \prec_{\mathsf{revlex}} x^{b'}$ as u is the rightmost entry where a' and b' differ and $a'_u > b'_u = 0$. So:

4.6.3. If a polytopal complex $\mathcal{C}$ has greedy representations, then $I_\mathcal{C} + \mathsf{in}_{\prec_{\mathsf{revlex}}}(I_U)$ is square-free.

Now making sure that we only count lattice points off a subcomplex $\mathcal{C}'$ is easy. Let $\mathcal{C}'$ be a subcomplex of $\mathcal{C}$. We consider its polytopal Stanley-Reisner ideal $I_{\mathcal{C}'}$ in $\mathbb{K}[x]$. If x^a is a monomial in $I_{\mathcal{C}'}$ but not in $I_{\mathcal{C}'} + \mathsf{in}_{\prec_{\mathsf{revlex}}}(I_U)$, then $\mathsf{supp}(a)$ is not contained in the vertex set of any face of $\mathcal{C}'$, but it is contained in the vertex set of some face F of $\mathcal{C}$. As $\mathcal{C}'$ is a subcomplex of $\mathcal{C}$, this means that Ua does not lie on $||a||_1 \bigcup \mathcal{C}'$. Conversely, any lattice point $z \in k\bigcup\mathcal{C} \setminus \bigcup\mathcal{C}'$ has a non-negative integral representation in terms of the vertices of some face of $\mathcal{C}$, but has no such representation in terms of the vertices of a face of $\mathcal{C}'$.

So the monomials x^a in $\mathbb{K}[x]$ of degree $\deg(x^a) = k$ such that $x^a \in I_{\mathcal{C}'}$ but $x^a \notin I_\mathcal{C} + \mathsf{in}_{\prec_{\mathsf{revlex}}}(I_U)$ are in bijection with the lattice points in $k\bigcup\mathcal{C} \setminus \bigcup\mathcal{C}'$. Note that $I_{\mathcal{C}'} + \mathsf{in}_{\prec_{\mathsf{revlex}}}(I_U)$ and $I_\mathcal{C} + \mathsf{in}_{\prec_{\mathsf{revlex}}}(I_U)$ are square-free and $I_{\mathcal{C}'} + \mathsf{in}_{\prec_{\mathsf{revlex}}}(I_U) \supset I_\mathcal{C} + \mathsf{in}_{\prec_{\mathsf{revlex}}}(I_U)$ and thus this pair of ideals can be viewed as a relative Stanley-Reisner ideal. This means we have just proved the following theorem.

4.6.4. Theorem.

Let $\mathcal{C}$ be a polytopal complex. If $\mathcal{C}$ has greedy representations, then for any subcomplex $\mathcal{C}' \subset \mathcal{C}$ there exists a relative Stanley-Reisner ideal $I_{\Delta,\Delta'}$ such that for all $0 < k \in \mathbb{N}$

$$\mathsf{L}_{\bigcup\mathcal{C}\setminus\bigcup\mathcal{C}'}(k) = \mathsf{H}_{I_{\Delta/\Delta'}}(k)$$

and this function is a polynomial.

We have obtained this theorem by looking at the proof we did in Section 4.3. The first step was to rephrase our definition of "stacking representations" in terms of the complex $\mathcal{C}$ given by the cube $[-1,1]^E$ intersected with the tension space on the one hand and the coordinate hyperplane arrangement on the other hand. This led to the concept of "greedy representations", which is applicable to any polytopal complex $\mathcal{C}$. Then by following our proof for Section 4.3 and translating the steps into the more abstract language of polytopal complexes and greedy representations, we showed that the Ehrhart function of a polytopal complex that has greedy representations is a Hilbert function.

Theorem 4.6.4 reads almost exactly like Theorem 4.4.4. This is no accident. In fact, these two theorems are identical, as is shown by the following theorem. To our knowledge neither the concept of greedy representations nor this theorem exist in the prior literature.

4.6.5. Theorem.

For a polytope P the following two statements are equivalent:

1. P is a compressed lattice polytope.
2. P has greedy representations.

Proof. *1. ⇒ 2.* Let $k \in \mathbb{N}$ and $z \in kP \cap \mathbb{Z}^d$. $\mathsf{minface}_P(\frac{1}{k}z)$ is non-empty, hence it has vertices, which by assumption are integral and so it contains a lattice point.

Let u be any such lattice point. We now consider the pulling triangulation $\mathsf{pull}(P, \prec)$ where $\prec$ is any total order on the lattice points that has u as its minimal element. Let σ be the inclusion-minimal face of $\mathsf{pull}(P, \prec)$ that contains $\frac{1}{k}z$, that is $\frac{1}{k}z \in \mathsf{relint}\,\sigma$. By 4.4.3 the lattice point u is a vertex of σ. Because σ is unimodular there is a non-negative integral representation a of z in terms of $\mathsf{vert}(\sigma)$ with $||a||_1 = k$. Finally $u \in \mathsf{supp}(a)$ as $\frac{1}{k}z$ lies in the relative interior of σ.

2. ⇒ 1. Let v be a vertex of P. By assumption v is rational, so let $k \in \mathbb{N}$ such that $kv \in \mathbb{Z}^d$. Then $kv \in kP$ and by assumption $\mathsf{minface}_P(\frac{1}{k}kv) = \{v\}$ contains a lattice point, which must be v itself. Hence P is a lattice polytope.

Let $\mathsf{pull}(P, \prec)$ be an arbitrary pulling triangulation of P and σ be one of its faces. We have to show that σ is unimodular. It suffices to show that for any $k \in \mathbb{N}$ and any $z \in \mathsf{relint}\,(k\sigma) \cap \mathbb{Z}^d$, there is a non-negative integral representation of z in terms of $\mathsf{vert}(\sigma)$. We proceed by induction on the dimension of σ. If $\mathsf{dim}\,(\sigma) = 0$, then σ is just a single lattice point and we are done. So let $\dim(\sigma) > 0$. Let F be the inclusion-minimal facet of P that contains σ and let u be the $\prec$-minimal lattice point in F. By 4.4.3 we know that u is a vertex of both σ and F. As P has greedy representations, we know that z has a non-negative integral representation a in terms of $\mathsf{vert}(F)$ with $||a||_1 = k$ and $u \in \mathsf{supp}(a)$.

Before we continue we prove the following claim: If $x \in l\sigma' \in \mathsf{pull}(P, \prec)$ such that $\sigma' \subset F$ and σ' is $\dim F$-dimensional, then $x + u \in (l+1)\sigma'$. As $x \in l\sigma'$ there is a unique non-negative representation a of x in terms of $\mathsf{vert}(\sigma')$ with $||x||_1 = l$. Then $a + e_u$ is a non-negative representation of $x + u$ in terms of $\mathsf{vert}(\sigma')$ with $||a + e_u||_1 = l + 1$. So $x + u \in (l+1)\sigma'$.

Now back to the proof proper. Consider $z - u \in (k-1)F$. We claim that moreover $z - u \in (k-1)\sigma$. For suppose $z - u$ were in $(k-1)\sigma'$ for a different $\dim F$-dimensional simplex σ' of $\mathsf{pull}(P, \prec)$, then by the above observation $z \in k\sigma'$, which is a contradiction to $z \in \mathsf{relint}(k\sigma)$. We keep subtracting u until we reach a $z - ju$ in the boundary of $(k-j)\sigma$. By induction on the dimension we know that $z - ju$ has a non-negative integral representation a' in terms of $\mathsf{vert}(\sigma)$ with $||a||_1 = k - j$. Then $a := a' + je_u$ is the desired representation of z. □

So Theorems 4.4.4 and 4.6.4 are the same theorem and in Sections 4.4 and 4.6 we have seen two proofs of this result that, at first glance, look very different. However, a closer look reveals that the main difference between the two proofs is this: We start out with a polytopal complex $\mathcal{C}$. In the first proof we first triangulate $\mathcal{C}$ and then pass to the corresponding Stanley-Reisner ideal. In the second proof we first pass to the corresponding polytopal Stanley-Reisner ideal and then add the initial ideal under the reverse lexicographic term order. There are other differences of course. In the first proof, for example, we delegate most of the work to Theorems 4.4.1 and 4.4.2 which we cite from [Sta96], while in the second proof we do all the work by hand.[11] Nonetheless, the crucial difference is "where" we do the real work: On the level of the polytopal complex or on the level of the corresponding ideal? This decision determines the language in which the proof is phrased and correspondingly the language in which the properties "being compressed" and "having greedy representations", that we require of $\mathcal{C}$, are formulated. We have just seen that these two properties are equivalent. And again, this is no accident. The explanation is given by the following fact, which we paraphrase from Sturmfels' book [Stu96, Chapter 8].

> Let $\prec$ be a total order on the vertices of $\mathcal{C}$ and let $\prec_{\mathsf{revlex}}$ be the reverse lexicographic term order on $\mathbb{K}[x]$ determined by the corresponding total order on the variables in $\mathbb{K}[x]$. Then:
>
> 1. $I_{\mathsf{pull}(\mathcal{C},\prec)} = I_{\mathcal{C}} + \mathsf{in}_{\prec_{\mathsf{revlex}}}(I_U)$.
> 2. $\mathsf{pull}(\mathcal{C}, \prec)$ is unimodular if and only if $\mathsf{in}_{\prec_{\mathsf{revlex}}}(I_U)$ is square-free.

More on this fascinating connection can be found in the chapter on regular triangulations in Sturmfels' book [Stu96, Chapter 8] and in the article [Stu91].

[11] Similarly in the applications: When we apply Theorem 4.4.4 to the tension polynomial in Section 4.5 we use the result of Ohsugi and Hibi, while when we "apply" Theorem 4.6.4 to the tension polynomial in Section 4.3 we prove that the boundary of the tension polytope has greedy representations by hand.

4.7. Non-Square-Free Ideals and Non-Standard Gradings

The proof we have seen in the last section suggests that we might be able to do with weaker assumptions, if we go for weaker conclusions. At the same time, we might save ourselves a lot of effort during the proof. So in this section we pursue two questions: What if we do not require the ideals to be square-free? And what if we allow our polynomial ring to be equipped with a non-standard grading?

It turns out that the answers are simple. If we only require a standard grading and not that our ideals are square-free, we can do with normal lattice polytopes in our complex. If we require neither, we can do with arbitrary lattice polytopes.

A lattice polytope P is **normal** if for every $k \in \mathbb{N}$ every $z \in kP \cap \mathbb{Z}^d$ can be written as the sum of k points in $P \cap \mathbb{Z}^d$. Note that a compressed polytope is automatically normal. A **non-standard grading** of $\mathbb{K}[x_1, \ldots, x_n]$ is given by a function $\deg : \{x_1, \ldots, x_n\} \to \mathbb{Z}_{\geq 0}$ that defines the degree of each variable. In the standard grading deg is constant 1. We now want to consider grading where each variable x_i may have a non-negative degree $\deg(x_i)$. The degree $\deg(x^a)$ of a monomial x^a is then defined as

$$\deg(x^a) = \deg(\prod_{i=1}^{n} x_i^{a_i}) = \sum_{i=1}^{n} a_i \deg(x_i).$$

Note that with this definition still $\deg(x^a x^b) = \deg(x^a) + \deg(x^b)$. Given two monomial ideals $I_2 \supset I_1$ in a polynomial ring $\mathbb{K}[x]$ equipped with some grading deg, we denote by H_{I_2/I_1} the Hilbert function that assigns to any $k \in \mathbb{N}$ the number of monomials in I_2 but not in I_1 that are of degree k with respect to the given grading.

With these two notions we can generalize Theorem 4.4.4, to include cases where the polytopal complex in question satisfies weaker conditions. The conclusions we obtain in these cases will not be as strong, however.

4.7.1. Theorem.

Let $\mathcal{C}$ be a polytopal complex in which all faces are lattice polytopes. Then for any subcomplex $\mathcal{C}' \subset \mathcal{C}$ there exist monomial ideals $I_2 \supset I_1$ in a polynomial ring $\mathbb{K}[x]$ equipped with some grading such that for all $0 < k \in \mathbb{N}$

$$\mathsf{L}_{\bigcup \mathcal{C} \setminus \bigcup \mathcal{C}'}(k) = \mathsf{H}_{I_2/I_1}(k)$$

and this function is a polynomial. If the faces of $\mathcal{C}$ are normal, the grading of $\mathbb{K}[x]$ can be chosen to be the standard grading. If the faces of $\mathcal{C}$ are compressed, then moreover the ideals I_1 and I_2 can be chosen to be square-free.

The case where the faces of $\mathcal{C}$ are compressed, the grading is standard and the ideals are square-free is just Theorem 4.4.4 of which we have already seen two proofs. We are not going to prove this again. The case where the faces are normal and the grading is standard is joint work with Aaron Dall, as is the rest of this chapter. The only exception is the case of this theorem where we pose no additional constraints on the faces of $\mathcal{C}$, the grading may be non-standard and the ideals may be non-square-free - this part is joint work with Raman Sanyal.

Proof. As in the previous section we assume without loss of generality that for every lattice point z there is at most one integer k such that $z \in k \bigcup \mathcal{C}$.

For any polytope $P \in \mathcal{C}$, there is a finite set of lattice points $\mathcal{U}_P$ in $\bigcup_{k=1}^{\infty} kP$ such that every lattice point in $\bigcup_{k=1}^{\infty} kP$ has a non-negative integral representation in terms of points in $\mathcal{U}_P$. (This is saying that every rational polyhedral cone has a Hilbert basis, see [Sch86, Theorem 16.4].) If P is normal, then $\mathcal{U}_P$ can be chosen such that $\mathcal{U}_P \subset P$.

Now consider the polynomial ring $\mathbb{K}[x_{u,P} : P \in \mathcal{C}, u \in \mathcal{U}_P]$ which is graded such that $\deg(x_{u,P})$ is the unique integer with $u \in kP$. In other words the variables of $\mathbb{K}[x]$ are indexed by *both* a lattice point u and a $P \in \mathcal{C}$ such that $u \in \deg(x_{u,P})P$. Note that if P is normal, then $\deg(x_u) = 1$ for all variables x_u. We fix an arbitrary total order on the variables $x_{u,P}$ and denote the associated reverse lexicographic term order by $\prec_{\mathsf{revlex}}$. We let U be the matrix that has the vectors u as columns such that the columns indexed by pairs (u, P); this means that a given column may appear more than once. Let n be the number of columns of U and let d be the number of rows. Again we consider the maps $\pi : a \mapsto Ua$ and $\hat{\pi} : x^a \mapsto z^{Ua}$. Let

$$\begin{aligned} I_{\mathcal{C}} &:= \langle x^a \mid \text{there is no } Q \in \mathcal{C} \text{ such that } Q = P \text{ for all pairs } (u,P) \in \mathsf{supp}(a) \rangle, \\ I_1 &:= I_{\mathcal{C}} + \langle x^b | x^b \notin I_{\mathcal{C}}, \exists x^a \notin I_{\mathcal{C}} : Ua = Ub \wedge x^a \prec_{\mathsf{revlex}} x^b \rangle, \\ I_2 &:= \langle x^a \notin I_1 \mid \text{there is no } Q \in \mathcal{C}' \text{ such that } u \in Q \text{ for all pairs } (u,P) \in \mathsf{supp}(a) \rangle. \end{aligned}$$

We claim that these two ideals I_1 and I_2 have the desired properties. The construction is almost the same as before. The difference is that a single lattice point u may now correspond to multiple variables $x_{u,P}$ and $x_{u,Q}$. This change makes some of the following arguments easier. Also this leads to two different ways of defining the "polytopal Stanley-Reisner ideal"; we make use of both approaches in the definitions of $I_{\mathcal{C}}$ and I_2, respectively. The first way has the advantage that it guarantees that all variables in monomials *not* in $I_{\mathcal{C}}$ are labeled with the same polytope $P \in \mathcal{C}$. The second way guarantees that variables in monomials *in* I_2 are not labeled with lattice points that all lie in a common face of $\mathcal{C}'$. We call an $a \in \mathbb{Z}^n_{\geq 0}$ valid if $x^a \notin I_1$ but $x^a \in I_2$.

We are now going to show that $\hat{\pi}$ gives a bijection between the monomials $x^a \in I_2 \setminus I_1$ of degree k and the Laurent monomials z^v with $v \in \deg(x^a)(\bigcup \mathcal{C} \setminus \bigcup \mathcal{C}') \cap \mathbb{Z}^d$. We proceed in several steps.

If $x^a \in I_2 \setminus I_1$, then $Ua \in \deg(x^a) \bigcup \mathcal{C} \cap \mathbb{Z}^d$. All variables appearing in x^a are labeled with a common polytope $P \in \mathcal{C}$. So $x^a = \prod_{u \in \mathcal{U}_P} x_{u,P}^{a_{u,P}}$. But then

$$Ua = \sum_{u \in \mathcal{U}_P} a_{u,P} u \in \sum_{u \in \mathcal{U}_P} a_{u,P} \deg(x_{u,P}) P = \deg(x^a) P.$$

$P \notin \mathcal{C}'$ by construction of I_2. Suppose $Ua \in \deg(x^a) Q$ for some $Q \in \mathcal{C}'$. Because $\mathcal{C}'$ is a subcomplex of the polytopal complex $\mathcal{C}$, we know that Q is a proper face of P and all the lattice points u with $x_{u,P} \in \mathsf{supp}(a)$ are also lattice points in Q. But then by construction of I_2, we know $x^a \notin I_2$, a contradiction.

$\hat{\pi}$ is surjective. Let $v \in k(\bigcup \mathcal{C} \setminus \bigcup \mathcal{C}') \cap \mathbb{Z}^d$ for some k. Then there is a polytope $P \in \mathcal{C} \setminus \mathcal{C}'$ such that $v \in \mathsf{relint}\,(kP)$. By construction of $\mathcal{U}_P$, v has a non-negative, integral representation in terms of elements of $\mathcal{U}_P$: $v = \sum_{u \in \mathcal{U}_P} b_u u$. Define $b \in \mathbb{Z}^n_{\geq 0}$ by $b_{u,P} := b_u$ if $u \in \mathcal{U}_P$ and $b_{w,Q} = 0$ if $Q \neq P$. So by construction $x^b \notin I_{\mathcal{C}}$ and $Ub = v$. Consider the $\prec_{\mathsf{revlex}}$-minimal monomial $x^a \notin I_{\mathcal{C}}$ with $Ua = Ub$. For this monomial we have $x^a \notin I_1$. Moreover, as $Ua = v \in \mathsf{relint}\,(kP)$ we have $x^a \in I_2$. Finally we have to check that $\deg(x^a) = k$. We have already seen that $\deg(x^a)$ is an integer with the property $Ua \in \deg(x^a)P$. However by our assumption at the beginning there is at most one integer k' such that $v = Ua \in k'P$. Thus $\deg(x^a) = k$.

$\hat{\pi}$ is injective. By definition of I_1, for every $v \in k(\bigcup \mathcal{C} \setminus \bigcup \mathcal{C}') \cap \mathbb{Z}^d$ there is at most one monomial $x^a \notin I_1$ such that $Ua = v$.

□

4.8. *The Structure of the Tension Polytope and Tension Complex*

In Theorem 4.5.4 we captured the integral tension polynomial as the Hilbert function of a relative Stanley-Reisner ideal $I_{\Delta/\Delta'}$. In our construction, the simplicial complexes $\Delta \supset \Delta'$ are given as geometric simplicial complexes. To extract more information about the integral tension polynomial from our construction we need to arrive at a better understanding of Δ and Δ'. So in this section we take a first few steps in this direction: We give a combinatorial characterization of the face lattice of the tension polytope $\bigcup \Delta$ as given by Theorem 4.5.4 and show that the variant of the tension complex Δ' given by Theorem 4.3.9 is homeomorphic as a topological space to the coloring complex defined by Steingrímsson.

The **tension polytope** T of a graph $G = (V, E)$ is the tension space of G intersected with the cube $[-1, 1]^E$, i.e. T consists of all $\mathbb{R}$-tensions t with $-1 \leq t \leq 1$. This is precisely $\bigcup \Delta$ where Δ is the complex given by the construction used in Theorem 4.5.4. A **skeleton** is a *partial* function $s : E \to \{+, -\}$ such that there exists a tension $t \in T$ with

$$s^{-1}(+) = \{e \in E | t(e) = +1\} \text{ and } s^{-1}(-) = \{e \in E | t(e) = -1\}. \tag{4.10}$$

Given a tension t we denote the partial function s defined by (4.10) by $s(t)$. A ***d*-skeleton** is a skeleton s such that

$$c(G[\mathsf{supp}(s)]) = c(G) + d.$$

Recall that c denotes the number of components. We equip the set of skeletons of G with the partial order given by the reverse inclusion. That is $s_1 \leq s_2$ iff $s_1 \supset s_2$, i.e. iff $\mathsf{supp}(s_1) \supset \mathsf{supp}(s_2)$ and $s_1|_{\mathsf{supp}(s_2)} = s_2$.

4.8.1. Theorem.

The poset of skeletons of G is the face lattice of T. The d-skeletons correspond to the d-faces of G.

Proof. The $\mathbb{R}$-tensions in the relative interior of a face of the tension polytope are precisely those tensions $t \in T$ such that a certain subset of the inequalities $-1 \leq t \leq 1$ hold at equality. That is, any non-empty face F of T is of the form

$$F = T \cap \{t(e) = -1 \mid e \in S_-\} \cap \{t(e) = +1 \mid e \in S_+\} \tag{4.11}$$

for some disjoint sets $S_-, S_+ \subset E$. Note that S_-, S_+ are not uniquely determined by F, but there are unique inclusion-maximal sets S_-, S_+ such that (4.11) holds. So we define $s(F)$ to be the partial function $s(F) = s : E \to \{+, -\}$ such that

$$s^{-1}(+) = \{e \in E \,|\, \forall t \in F : t(e) = +\} \text{ and } s^{-1}(-) = \{e \in E \,|\, \forall t \in F : t(e) = -\}.$$

Note that $s(F) = s(t)$ for any $t \in \mathsf{relint}\, F$ and $\mathsf{relint}\, F = \{t \in T | s(t) = s(F)\}$. If F' is a face of F then $s(F') \supset s(F)$. This already shows that the poset of skeletons of G, ordered by reverse inclusion, is the face lattice of T (minus the empty set).

We still have to show that $s(F)$ is a $\dim F$-skeleton. We proceed in several steps. Throughout we fix $t \in \mathsf{relint}\, F$.

$\mathsf{aff}(F) = \{x \textit{ tension} \,|\, s(x) \supset s(t)\}$. The set $\{x \text{ tension} \,|\, s(x) \supset s(t)\}$ is affine and contains F. We still have to show that any tension x with $s(x) \supset s(t)$ is contained in $\mathsf{aff}(F)$. If $\mathsf{supp}(s(t)) = E$, then $x = t$ and we are done. Otherwise we consider $x - t$. $\mathsf{zero}(x - t) \supset \mathsf{supp}(s(t))$. Let $\epsilon > 0$ such that $-1 < t + \epsilon(x - t) < 1$. Such an ϵ exists because $t \in \mathsf{relint}\, F$. Now $s(t + \epsilon(x - t)) \supset s(t)$ because $\mathsf{zero}(\epsilon(x - t)) \supset \mathsf{supp}(s(t))$. So we have $t + \epsilon(x - t) \in F$ and $t \in F$ and consequently $x \in \mathsf{aff}\{t + \epsilon(x - t), t\} \subset \mathsf{aff}(F)$.

$\mathsf{aff}(F) - t = \{t' \textit{ tension} | \mathsf{zero}(t') \supset \mathsf{supp}(s(t))\}$. If $t' \in \mathsf{aff}(F) - t$, then $\mathsf{zero}(t') \supset \mathsf{supp}(s(t))$ because $t' = x - t$ for some tension x with $s(x) \supset s(t)$. Conversely if t' is a tension with $\mathsf{zero}(t') \supset \mathsf{supp}(s(t))$, then $s(t' + t) \supset s(t)$ and so $t' \in \mathsf{aff}(F) - t$.

However we have already seen that for tensions fixing the weight of an edge to be zero amounts to the same thing as contracting the edge in the graph (cf. 3.8.5). So $\mathsf{aff}(F) - t$ *is isomorphic to the tension space of* $G/\mathsf{supp}(s(t))$.

The dimension of the tension space of $G/\mathsf{supp}(s(t))$ is the number of vertices of $G/\mathsf{supp}(s(t))$ minus the number of components of $G/\mathsf{supp}(s(t))$. The number of vertices of $G/\mathsf{supp}(s(t))$ is just the number of components of $G[\mathsf{supp}(s(t))]$, while the number of components of $G/\mathsf{supp}(s(t))$ is the number of components of G. So

$$\dim F = c(G[\mathsf{supp}(s(t))]) - c(G)$$

as desired. □

Let us now compare the tension complex and the coloring complex. The first difficulty in this task is that we have three slightly different results about the integral tension polynomial (4.3.9, 4.3.10, 4.5.4) that each give rise to a different notion of tension complex. Also Steingrímsson defines two different complexes: one is given by the Stanley-Reisner ideal I_2 from Section 4.2, while the other is what is referred to in the literature as "the coloring complex". Moreover there is the complex given by Theorem 4.5.11 to consider, which does not appear in Steingrímsson's work. Let us examine each of these in turn.

The simplicial complex $\mathcal{C}_1$ defined by the ideal I_2 from Theorem 4.3.9 is, speaking geometrically, the intersection of the coordinate hyperplanes with the *boundary* of the tension polytope.[12] The complex $\mathcal{C}_2$ given by I_2 in Corollary 4.3.10 can be obtained by constructing the pyramid over $\mathcal{C}_2$, i.e. by considering the complex whose facets are of the form $\mathrm{conv}\{0, F\}$ where F ranges over the facets of $\mathcal{C}_2$. Alternatively, $\mathcal{C}_2$ is the intersection of the coordinate hyperplanes with the entire tension polytope. Finally let $\mathcal{C}_3$ denote the complex Δ' from Theorem 4.5.4. $\mathcal{C}_3$ is the union of $\mathcal{C}_2$ and the boundary of the tension polytope.

The complex Δ_1 given by the ideal I_2 constructed in Section 4.2 can be described as follows. The graphic arrangement given by G consists of those hyperplanes $\{x \mid x_u = x_v\}$ in the braid arrangement for which $uv = e \in E$ is an edge of the graph. Δ_1 is now the intersection of the graphic arrangement with the unit cube. Each of the hyperplanes in the braid arrangement contains the diagonal $\mathrm{conv}\{0, 1\} = \{x \mid x_{v_1} = x_{v_2} = \ldots = x_{v_n}\}$ where $v_1, \ldots, v_n$ are the vertices of the graph G. So Steingrímsson went on to define the link of this diagonal to be the coloring complex Δ_2 as it is called in the literature. Conversely, the complex Δ_1 is the double pyramid over the coloring complex Δ_2. One alternative description is this: Consider the boundary of the cube as a polytopal complex and remove the vertices 0 and 1, as well as all faces containing either. Δ_2 is the intersection of what is left of the cube with the graphic arrangement. Another option is to obtain Δ_2 by intersecting the braid arrangement with the unit sphere of dimension $n-2$ in the hyperplane $\{x \mid \sum_{i=1}^{n} x_i = 0\}$ and then considering Δ_2 to be the subcomplex consisting of all faces that lie on a hyperplane of the graphic arrangement. See [HS08, Theorem 1]. Finally let Δ_3 be the complex Δ' given by Theorem 4.5.11. Δ_3 is the union of Δ_1 with all the faces of the unit cube where at least one coordinate is one, i.e. with all faces $\{x \in [0,1]^V \mid x_v = 1\}$ for $v \in V$. Note the difference in construction between $\mathcal{C}_3$ and Δ_3: while $\mathcal{C}_3$ is the union of $\mathcal{C}_2$ with a topological sphere, Δ_3 is the union of Δ_1 with a topological disk.

If we consider the constructions in the integral tension case and in the chromatic case side by side, we note that $\mathcal{C}_1$ and Δ_2 are most similar. $\mathcal{C}_2$ and Δ_1 also seem to be related, but $\mathcal{C}_2$ is $|V| - c - 1$-dimensional, while Δ_1 is $|V| - 1$-dimensional. If G is connected, however, then $\mathcal{C}_1$ and Δ_2 are both $|V| - 3$-dimensional. Indeed we have the following theorem.

4.8.2. Theorem.

The tension complex $\mathcal{C}_1$ and the coloring complex Δ_2 of a connected graph are homeomorphic as topological spaces.

On the one hand a close relation between the two complexes is to be expected, given the close relation between tensions and colorings, and indeed the homeomorphism is between the complexes is induced by something as simple as the incidence matrix. On the other hand this comes as something of a surprise as the tension complex determines the *integral* tension polynomial θ_G while the coloring complex determines the chromatic polynomial which is a multiple only of the *modular* tension polynomial $\bar{\theta}_G$ and not of θ_G.

[12]Of course we also have to triangulate appropriately. To make the presentation more concise we will not mention the necessary triangulation steps here.

Proof. We consider the linear map from $\mathbb{R}^E$ into $\mathbb{R}^V$ defined by $e_{uv} \mapsto e_v - e_u$ for any edge $uv \in E$. The matrix of this linear map is just the incidence matrix A of G. Recall that $\ker A$ is the flow space of G and $\ker A^\perp$ is the tension space of G. By 3.4.1 and the assumption that G is connected the image of A is the set of vectors $b \in \mathbb{R}^V$ with $\sum_{v\in V} b_v = 0$. So A gives a linear isomorphism between the tension space and the set $\{x \in \mathbb{R}^V \,|\, \sum_{v\in V} x_v = 0\}$. Another way of putting this would be to say that A maps a tension t to the corresponding coloring $c = At$ with the property that $\sum_{v\in V} c_v = 0$.

Now, for any edge $e \in E$ from u to v

$$A\{x \in \ker A^\perp \mid x_e = 0\} = \{x \in \mathbb{R}^V \mid x_u = x_v, \sum_{v\in V} x_v = 0\}.$$

since if x is a tension and $x_e = 0$ then $x_u = x_v$ and

$$\dim\{x \in \ker A^\perp \mid x_e = 0\} = |V| - 2 = \dim\{x \in \mathbb{R}^V \mid x_u = x_v, \sum_{v\in V} x_v = 0\}.$$

So the linear map φ gives a homeomorphism between the coordinate hyperplane arrangement in tension space and the graphic arrangement in $\{x \in \mathbb{R}^V \,|\, \sum_{v\in V} x_v = 0\}$. As $\mathcal{C}_1$ is homeomorphic to the intersection of the unit sphere in tension space with the coordinate hyperplane arrangement and Δ_2 is homeomorphic to the intersection of the unit sphere in $\{x \in \mathbb{R}^V \,|\, \sum_{v\in V} x_v = 0\}$ with the braid arrangement, this completes the proof. □

The proof suggests that it may even be possible to describe the combinatorial relationship between the tension and coloring complexes. It may be interesting to investigate this connection further, as this might shed more light on the relationship between the integral and modular tension polynomials.[13]

4.9. *Bounds on the Coefficients*

This section is about the problem of obtaining bounds for the coefficients of the five counting polynomials. We do not attempt to give a comprehensive treatment of the subject, we only make a first few observations: We characterize polynomials that are Hilbert functions of relative Stanley-Reisner ideals and use this characterization to derive some easy bounds on the coefficients of our five counting polynomials. We conclude by suggesting ways in which these bounds might be improved, some of which seem quite feasible.

$\binom{k-1}{d}$ is a polynomial of degree d in k. The polynomials $\binom{k-1}{i}$ for $0 \le i \le d$ form a basis of the $\mathbb{K}$-vector space of all polynomials in $\mathbb{K}[k]$ of degree at most d and the polynomials $\binom{k-1}{i}$ for $0 \le i$ form a basis of $\mathbb{K}[k]$ when seen as a $\mathbb{K}$-vector space. By Theorem 4.4.2 the coefficients of the Hilbert function of a relative Stanley-Reisner ideal expressed with respect to this basis must be non-negative and integral. It turns out that this characterizes which polynomials appear as Hilbert functions of relative Stanley-Reisner ideals.

[13] Recall that the modular tension polynomial is a divisor of the chromatic polynomial.

4.9.1. Theorem.

A polynomial $f(k) = \sum_{i=0}^{d} f_i \binom{k-1}{i}$ is the Hilbert function of some relative Stanley-Reisner ideal $I_{\Delta/\Delta'}$ if and only if $0 \leq f_i \in \mathbb{N}$ for all $0 \leq i \leq d$.

Proof. We have already seen that the coefficients of $\mathsf{H}_{I_{\Delta/\Delta'}}(k)$ with respect to the basis $\binom{k-1}{d}$, $d \in \mathbb{Z}_{\geq 0}$ are necessarily non-negative integers. To see that this is also sufficient, let $f(k) = \sum_{i=0}^{d} f_i \binom{k-1}{i}$ with $0 \leq f_i \in \mathbb{N}$ for all $0 \leq i \leq d$. For $0 \leq i \leq d$ and $1 \leq j \leq f_i$ let σ_j^i denote a closed unimodular lattice simplex of dimension i in $\mathbb{R}^d$ such that the σ_j^i are pairwise disjoint. Let Δ denote the (disjoint) union of all these σ_j^i and define Δ' to be the union of the respective boundaries $\partial\sigma_j^i$. Then the set $\bigcup \Delta \setminus \bigcup \Delta'$ is the disjoint union of f_d relatively open unimodular lattice simplices of dimension d, f_{d-1} relatively open unimodular lattice simplices of dimension $d-1$ and so on. Consequently $\mathsf{H}_{I_{\Delta/\Delta'}}(k) = \mathsf{L}_{\bigcup \Delta \setminus \Delta'}(k) = f(k)$ as desired. □

This immediately implies that all the counting functions we considered have non-negative integral coefficients with respect to this basis.

4.9.2. Theorem.

The k-flow and $\mathbb{Z}_k$-flow polynomials, the k-tension and $\mathbb{Z}_k$-tension polynomials and the chromatic polynomial of a graph have non-negative integer coefficients with respect to the basis $\{\binom{k-1}{d} | 0 \leq d \in \mathbb{Z}\}$ of $\mathbb{K}[k]$.

The bounds on the coefficients of the chromatic polynomial are much weaker then the bounds given in [HS08] which by 3.2.6 also apply to the modular tension polynomial. We have been unable to find equivalent or stronger bounds on the coefficients of the modular flow polynomial and the integral flow and tension polynomials in the prior literature. To obtain more information about the coefficients we would have to take more combinatorial and topological information about the complexes Δ and Δ' into account.

Theorem 4.9.1 is reminiscent of Macaulay's characterization of the Hilbert polynomials of graded algebras.

4.9.3. Theorem. (Macaulay, cf. [Bre98, Theorem 2.5])

A polynomial $f \in \mathbb{Q}[k]$ of degree d with $f(\mathbb{Z}) \subset \mathbb{Z}$ for which there exist integers $m_0, \ldots, m_1$ such that

$$f(k) = \sum_{i=0}^{d} \binom{k+i}{i+1} - \binom{k+i-m_i}{i+1}$$

is the Hilbert polynomial of a standard graded algebra if and only if

$$m_0 \geq m_1 \geq \ldots \geq m_d \geq 0.$$

As we already mentioned, the Hilbert functions of relative Stanley-Reisner rings are not in general Hilbert functions or Hilbert polynomials of standard graded algebras, so the relationship between Macaulay's Theorem and Theorem 4.9.1 is not immediately clear.

We conclude this thesis with some informal comments on the issue of obtaining bounds on the coefficients and directions for future research.

The strongest bounds on the coefficients of the chromatic polynomial are given by Hersh and Swartz [HS08]. These bounds are derived from the result that the coloring complex has a convex ear decomposition, a notion which was introduced by Chari [Cha97]. We refer to these two articles for a definition of this concept. We have nothing new to add to these results, even though we obtain the chromatic polynomial itself as a Hilbert function and not only the shifted chromatic polynomial. The reason is that the complex Δ' given in Theorem 4.5.11 is of a different structure. In particular it does not have a convex ear decomposition - the reason is that we are working with a half-open cube. So while new bounds might be derived on the basis of our alternative coloring complex Δ', a different approach would have to be used.

In the case of the integral flow and tension polynomials, the situation is different. Here our geometric setup strongly suggests that the complexes Δ' given by Theorems 4.5.6 and 4.5.4, respectively, have a convex ear decomposition. Especially as we use regular triangulations of the polytopal complexes we start out with.

The bad news, however, is that even this result would not lead to useful bounds on the coefficients. The basic setup in all five cases is such that the counting polynomial in question, let us call it f, is of the form $f = \mathsf{L}_\Delta - \mathsf{L}_{\Delta'}$, i.e. it is the difference of two Ehrhart polynomials. The nice thing about the coloring case is that L_Δ is particularly easy to describe: It is simply the Ehrhart function of the cube. Thus bounds on the h-vector of Δ' and thus on the coefficients of $\mathsf{L}_{\Delta'}$ translate directly into bounds on the coefficients of f. But in the case of integral flows and tension neither L_Δ nor $\mathsf{L}_{\Delta'}$ have a simple form. One would have to obtain bounds on the coefficients of both L_Δ and $\mathsf{L}_{\Delta'}$, and get only very weak bounds on their difference in return.

In the modular case, on the other hand, we are in good shape. There L_Δ is again the Ehrhart polynomial of a cube. And using the inside-out polytope construction in both the modular flow and tension case, Δ' is the union of the boundary of this cube and a hyperplane arrangement intersected with the cube. So again the geometric setup suggests that these complexes have a convex ear decomposition. Proving this statement would then translate into new bounds on the coefficients, at least for the modular flow polynomial.

So we end this thesis with the spirited remark:

There is a lot of work to do!

Bibliography

[Aig07] Martin Aigner, *A course in enumeration*, Graduate Texts in Mathematics, vol. 238, Springer-Verlag, 2007.

[Bar94] Alexander Barvinok, *A polynomial time algorithm for counting integral points in polyhedra when the dimension is fixed*, Mathematics of Operations Research **19** (1994), no. 4, 769–779.

[Bar08] ______, *Integer points in polyhedra*, Zürich Lectures in Advanced Mathematics, European Mathematical Society, 2008.

[BB] Matthias Beck and Eric Babson, *Minimal-distance chromatic and modular flow polynomials*, in preparation.

[BCJ01] G.B. Byrnes, G. Cairns, and B. Jessup, *Left-overs from the ham-sandwich theorem*, American Mathematical Monthly **108** (2001), 246–249.

[BGMM99] Hélymar Balza-Gomez, Jean-Michel Moreau, and Dominique Michelucci, *Convex hull of grid points below a line or a convex curve*, Proceedings of the 8th International Conference on Discrete Geometry for Computer Imagery (London), Springer-Verlag, 1999, pp. 361–374.

[BHJ08] Imre Bárány, Alfredo Hubard, and Jesús Jerónimo, *Slicing convex sets and measures by a hyperplane*, Discrete & Computational Geometry **39** (2008), no. 1-3, 67.

[BHM08] Matthias Beck, Christian Haase, and Asia R. Matthews, *Dedekind-Carlitz polynomials as lattice-point enumerators in rational polyhedra*, Mathematische Annalen **341** (2008), no. 4, 945–961.

[Bir12] George D. Birkhoff, *A determinant formula for the number of ways of coloring a map*, The Annals of Mathematics **14** (1912), no. 1/4, 42–46.

[BKS00] S. Bespamyatnikh, D. Kirkpatrick, and J. Snoeyink, *Generalizing ham sandwich cuts to equitable subdivisions*, Discrete Comput. Geom. **24** (2000), 605–622.

[BM01] I. Bárány and J. Matoušek, *Simultaneous partitions of measures by k-fans*, Discrete Comput. Geom. **25** (2001), 317–334.

[BO92] Thomas Brylawski and James Oxley, *The Tutte polynomial and its applications*, Matroid applications, vol. 40, Cambridge Univ. Press, Cambridge, 1992, pp. 123–225.

[BR07] Matthias Beck and Sinai Robins, *Computing the continuous discretely*, Undergraduate Texts in Mathematics, Springer, New York, 2007.

[Bre98] Francesco Brenti, *Hilbert polynomials in combinatorics*, Journal of Algebraic Combinatorics **7** (1998), 127–156.

[Bre09] Felix Breuer, *Uneven splitting of ham sandwiches*, Discrete and Computational Geometry (2009).

[Bro93] Tom C. Brown, *Descriptions of the characteristic sequence of an irrational*, Canadian Mathematical Bulletin **36** (1993), no. 1, 15–21.

[BZ04] W.A. Beyer and A. Zardecki, *The early history of the ham sandwich theorem*, American Mathematical Monthly **111** (2004), 58–61.

[BZ06a] Matthias Beck and Thomas Zaslavsky, *Inside-out polytopes*, Adv. Math. **205** (2006), no. 1, 134–162.

[BZ06b] ______, *The number of nowhere-zero flows on graphs and signed graphs*, J. Combin. Theory Ser. B **96** (2006), no. 6, 901–918.

[Cha97] Manoj K. Chari, *Two decompositions in topological combinatorics with applications to matroid complexes*, Transactions of the American Mathematical Society **349** (1997), no. 10, 3925–3943.

[Che07] Beifang Chen, *Orientations, lattice polytopes, and group arrangements I: Chromatic and tension polynomials of graphs*, arXiv:0706.3273v1, 2007.

[Dal08] Aaron Dall, *The flow and tension complexes*, Master's thesis, San Francisco State University, 2008.

[DEBG$^+$06] A. Dumitrescu, A. Ebbers-Baumann, A. Grüne, R. Klein, and G. Rote, *On the geometric dilation of closed curves, graphs and point sets*, Computational Geometry, Theory and Applications **36** (2006), 16–38.

[Die05] Reinhard Diestel, *Graph theory*, third ed., Graduate Texts in Mathematics, vol. 173, Springer-Verlag, 2005.

[DL05] Jesús A. De Loera, *The many aspects of counting lattice points in polytopes*, Mathematische Semesterberichte **52** (2005), no. 2, 175–195.

[DLHTY04] Jesús A. De Loera, Raymond Hemmecke, Jeremiah Tauzer, and Ruriko Yoshida, *Effective lattice point counting in rational convex polytopes*, Journal of Symbolic Computation **38** (2004), no. 4, 1273–1302.

[Dol92] V.L. Dol'nikov, *A generalization of the ham sandwich theorem*, Math. Notes **52** (1992), 771–779.

[Ehr77] E. Ehrhart, *Polynômes arithmétiques et méthode des polyèdres en combinatoire*, Birkhäuser Verlag, Basel, 1977, International Series of Numerical Mathematics, Vol. 35.

[Eis94] David Eisenbud, *Commutative algebra with a view toward algebraic geometry*, Graduate Texts in Mathematics, vol. 150, Springer-Verlag, 1994.

[EMM08] Joanna Ellis-Monaghan and Criel Merino, *Graph polynomials and their applications I: The Tutte polynomial*, arXiv:0803.3079v2, 2008.

[FMT78] Aviezri S. Fraenkel, Mordechai Mushkin, and Uzi Tassa, *Determination of* $[n\theta]$ *by its sequences of differences*, Canadian Mathematical Bulletin **21** (1978), 441–446.

[Fra05] Aviezri S. Fraenkel, *Complementary balanced words over an arbitrary finite alphabet*, http://www.wisdom.weizmann.ac.il/~fraenkel/ (2005).

[GH05] R. Guàrdia and F. Hurtado., *On the equipartition of plane convex bodies and convex polygons*, J. Geom. **83** (2005), 32–45.

[Gio07] Emeric Gioan, *Enumerating degree sequences in digraphs and a cycle-cocycle reversing system*, Eur. J. Comb. **28** (2007), no. 4, 1351–1366.

[GJ79] Michael R. Garey and David S. Johnson, *Computers and intractability: A guide to the theory of NP-completeness*, W. H. Freeman, 1979.

[GLL78] Ronald L. Graham, Shen Lin, and Chio-Shih Lin, *Spectra of numbers*, Mathematics Magazine **51** (1978), no. 3, 174–176.

[GR04] Chris Godsil and Gordon Royle, *Algebraic graph theory*, Graduate Texts in Mathematics, vol. 207, Springer-Verlag, 2004.

[Grü60] Branko Grünbaum, *Partitions of mass-distributions and convex bodies by hyperplanes*, Pacific Journal of Mathematics **10** (1960), 1257–1261.

[GZ83] Curtis Greene and Thomas Zaslavsky, *On the interpretation of Whitney numbers through arrangements of hyperplanes, zonotopes, non-Radon partitions, and orientations of graphs*, Trans. Amer. Math. Soc. **280** (1983), no. 1, 97–126. MR MR712251 (84k:05032)

[Hat02] Allen Hatcher, *Algebraic topology*, Cambridge University Press, 2002.

[HRW98] Jürgen Herzog, Vic Reiner, and Volkmar Welker, *The Koszul property in affine semigroup rings*, Pacific Journal of Mathematics **186** (1998), no. 1, 39–65.

[HS08] Patricia Hersh and Ed Swartz, *Coloring complexes and arrangements*, Journal of Algebraic Combinatorics **27** (2008), no. 2, 205–214.

[Hul07] A Hultman, *Link complexes of subspace arrangements*, European Journal of Combinatorics **28** (2007), no. 3, 781–790.

[Ist81] V.I. Istrăţescu, *Fixed point theory*, Mathematics and Its Applications, vol. 7, D. Reidel Publishing Company, 1981.

[Jon05] Jakob Jonsson, *The topology of the coloring complex*, Journal of Algebraic Combinatorics **21** (2005), 311–329.

[Kin08] J Kincses, *The topological type of the α-sections of convex sets*, Advances in Mathematics **217** (2008), no. 5, 2159–2169.

[KNA94] Naoyosho Kanamaru, Takao Nishizeki, and Tetsuo Asano, *Efficient enumeration of grid points in a convex polygon and its application to integer programming*, International Journal of Computational Geometry and Applications **4** (1994), 69–85.

[Koc02] Martin Kochol, *Polynomials associated with nowhere-zero flows*, J. Combin. Theory Ser. B **84** (2002), no. 2, 260–269.

[KRS99] W. Kook, V. Reiner, and D. Stanton, *A convolution formula for the Tutte polynomial*, J. Combin. Theory Ser. B **76** (1999), no. 2, 297–300.

[Kul97] W. Kulpa, *The Poincaré-Miranda theorem*, American Mathematical Monthly **104** (1997), 545–550.

[LRS09] Jesús A. De Loera, Jörg Rambau, and Francisco Santos, *Triangulations: Structures and algorithms*, book in preparation, 2009.

[Mat02] J. Matoušek, *Lectures on discrete geometry*, Graduate Texts in Mathematics, vol. 212, Springer-Verlag, 2002.

[Mat08] ______, *Using the Borsuk-Ulam theorem (corrected 2nd printing)*, Springer-Verlag, 2008.

[Mau81] R.D. Mauldin, *The Scotish book*, Birkhäuser, 1981.

[Meg85] N. Megiddo, *Partitioning with two lines in the plane*, J. Algorithms **6** (1985), 430–433.

[Mir40] C. Miranda, *Un' osservazione su una teorema di brouwer*, Boll. Unione Math. Ital. **3** (1940), 527.

[MS05] Ezra Miller and Bernd Sturmfels, *Combinatorial commutative algebra*, Graduate Texts in Mathematics, vol. 227, Springer-Verlag, 2005.

[Noo81] P. Noordzij, *On* $(na) + (nb) = (nc) + (nd)$ *for all positive integers* n, Nieuw Archief voor Wiskunde, III. Ser. **29** (1981), 59–70.

[O'B02] Kevin O'Bryant, *A generating function technique for Beatty sequences and other step sequences*, Journal of Number Theory **94** (2002), no. 2, 299–319.

[OH01] Hidefumi Ohsugi and Takayuki Hibi, *Convex polytopes all of whose reverse lexicographic initial ideals are squarefree*, Proceedings of the American Mathematical Society **129** (2001), no. 9, 2541–2546.

[Poi86] H. Poincaré, *Sur les courbes défines par une équation différentielle iv*, J. Math. Pures Appl. **85** (1886), 151–217.

[PS90] Horacio Porta and Kenneth B. Stolarsky, *Half-silvered mirrors and Wythoff's game*, Canadian Mathematical Bulletin **33** (1990), 119–125.

[Rad46] R. Rado, *A theorem on general measure*, J. London Math. Soc. **21** (1946), 291–300.

[Ram96] E.A. Ramos, *Equipartition of mass distributions by hyperplanes*, Discrete Comput. Geom. **15** (1996), 147–167.

[Ree57] John E. Reeve, *On the volume of lattice polyhedra*, Proceedings of the London Mathematical Society, II. Ser. **7** (1957), 378–395.

[Rei99] Victor Reiner, *An interpretation for the Tutte polynomial*, European J. Combin. **20** (1999), no. 2, 149–161.

[Rez06] Bruce Reznick, *Clean lattice tetrahedra*, arXiv.org:math/0606227 (2006).

[Sca85] Herbert E. Scarf, *Integral polyhedra in three space*, Mathematics of Operations Research **10** (1985), 403–438.

[Sch86] Alexander Schrijver, *Theory of linear and integer programming*, Wiley-Interscience Series in Discrete Mathematics, John Wiley & Sons Ltd., Chichester, 1986, A Wiley-Interscience Publication.

[ST42] A.H. Stone and J.W. Tukey, *Generalized 'sandwich' theorems*, Duke Math. J. **9** (1942), 356–359.

[Sta73] Richard P. Stanley, *Acyclic orientations of graphs*, Discrete Math. **5** (1973), 171–178.

[Sta96] ______, *Combinatorics and commutative algebra second edition*, Progress in Mathematics, vol. 41, Birkhäuser, 1996.

[Ste38] H. Steinhaus, *A note on the ham sandwich theorem*, Mathesis Polska **11** (1938), 26–28.

[Ste01] Einar Steingrímsson, *The coloring ideal and coloring complex of a graph*, Journal of Algebraic Combinatorics **14** (2001), 73–84.

[Sto76] Kenneth B. Stolarsky, *Beatty sequences, continued fractions, and certain shift operators*, Canadian Mathematical Bulletin **19** (1976), 473–482.

[Sto90] I. Stojmenović, *Bisections and ham-sandwich cuts of convex polygons and polyhedra*, Proceedings of the Second Canadian Conference in Computational Geometry, 1990, pp. 291–295.

[Sto91] ______, *Bisections and ham-sandwich cuts of convex polygons and polyhedra*, Information Processing Letters **38** (1991), 15–21.

[Stu91] Bernd Sturmfels, *Gröbner bases of toric varieties*, Tôhoku Math. Journal **43** (1991), 249–261.

[Stu96] ———, *Gröbner bases and convex polytopes*, University Lecture Series, vol. 8, American Mathematical Society, 1996.

[Sul04] Seth Sullivant, *Compressed polytopes and statistical disclosure limitation*, arXiv:math/0412535v1, 2004.

[Tut54] W. T. Tutte, *A contribution to the theory of chromatic polynomials*, Canadian J. Math. **6** (1954), 80–91. MR MR0061366 (15,814c)

[Vre03] S.T. Vrećica, *Tverberg's conjecture*, Discrete Comput. Geom. **29** (2003), 505–510.

[VŽ01] S.T. Vrećica and R.T. Živaljević, *Conical equipartitions of mass distributions*, Discrete Comput. Geom. **25** (2001), 335–350.

[Wes00] Douglas B. West, *Introduction to graph theory (2nd edition)*, Prentice Hall, August 2000.

[Whi32a] Hassler Whitney, *The coloring of graphs*, The Annals of Mathematics **33** (1932), no. 4, 688–718.

[Whi32b] Hassler Whitney, *A logical expansion in mathematics*, Bull. Amer. Math. Soc. **38** (1932), 572–579.

[Whi64] George K. White, *Lattice tetrahedra*, Canadian Journal of Mathematics **16** (1964), 389–396.

[Zie97] Günter M. Ziegler, *Lectures on polytopes*, Graduate Texts in Mathematics, vol. 152, Springer-Verlag, 1997.

[Zin21] K. Zindler, *Über konvexe gebilde II*, Monatsh. Math. **31** (1921), 25–56.

[ŽV90] R.T. Živaljević and S.T. Vrećica, *An extension of the ham sandwich theorem*, Bull. London Math. Soc. **22** (1990), 183–186.