A Riemann-Roch Theory for Sublattices of the Root Lattice A_n, Graph Automorphisms and Counting Cycles in Graphs

Dissertation zur Erlangung des
Grades des Docktors der Naturwissenschaften (Dr. rar. nat)
der Naturwissenschaftlich-Technischen Fakultäten
der Universtät des Saarlandes

voglegt von

Madhusudan Manjunath,
Saarbrücken.

12. März 2012

Tag des Kolloquiums:	20. Dezember 2011
Dekan:	Prof. Dr. Holger Hermanns
Prüfungsausschuss:	
Vorsitzender:	Prof. Dr. Markus Bläser
Gutachter:	Prof. Dr. Kurt Mehlhorn
	Prof. Dr. Christian Haase
	Dr. Omid Amini
Akademischer Mitarbeiter:	Dr. Saurabh Ray

To my parents and my sister.

Acknowledgments: I am extremely grateful to Omid Amini for firstly introducing me to the beautiful and rich topic of Riemann-Roch theory in discrete mathematics, for his collaboration on the topic and his constant help and advice on all academic matters, including his enomormous effort into improving the quality of exposition in this thesis. I am extremely grateful to Kurt Mehlhorn for the freedom he has provided me to pursue my interests, for his interest on my development as a researcher and for his collaboration. Vikram Sharma has been constant source of support and I specially thank him for inviting me to visit the Institute of Mathematical Sciences. Special thanks to Joachim Giesen for the mentoring and guidance that he provided me when I was a beginning graduate student. Thanks to Konstantinos Panagiotou and He Sun for the collaboration on which the last chapter of the thesis is based. I would like to thank Raghavendra Rao for his helpful comments on a first draft of the thesis. I am extremely grateful to Bernd Sturmfels for the generous advice and the directions he has been providing me. I have greatly benefitted with my contact with Frank-Olaf Schreyer and I specially thank him for always willing to help me out in learning algebraic geometry. I wish to thank Hannah Markwig for informing me about the seminar on tropical geometry and algebraic curves and for giving me the opportunity to participate and speak in the seminar. Special thanks to Khaled Elbasionni and Amr Elmasry for the stimulating discussions we had on pretty much all aspects of life. Last but definitely not the least, I am extremely indebted to my parents, my sister and friends for all the moral support they have provided me through these years.

Abstract

This thesis consists of two independent parts. In the first part of the thesis, we develop a Riemann-Roch theory for sublattices of the root lattice A_n extending the work of Baker and Norine (Advances in Mathematics, 215(**2**): 766-788, 2007) and study questions that arise from this theory. Our theory is based on the study of critical points of a certain simplicial distance function on a lattice and establishes connections between the Riemann-Roch theory and the Voronoi diagrams of lattices under certain simplicial distance functions. In particular, we provide a new geometric approach for the study of the Laplacian of graphs. As a consequence, we obtain a geometric proof of the Riemann-Roch theorem for graphs and generalise the result to other sub-lattices of A_n. Furthermore, we use the geometric approach to study the problem of computing the rank of a divisor on a finite multigraph G to obtain an algorithm that runs in polynomial time for a fixed number of vertices, in particular with running time $2^{O(n \log n)}\mathrm{poly}(\mathrm{size}(G))$ where n is the number of vertices of G. Motivated by this theory, we study a dimensionality reduction approach to the graph automorphism problem and we also obtain an algorithm for the related problem of counting automorphisms of graphs that is based on exponential sums.

In the second part of the thesis, we develop an approach, based on complex-valued hash functions, to count cycles in graphs in the data streaming model. Our algorithm is based on the idea of computing instances of complex-valued random variables over the given stream and improves drastically upon the naïve sampling algorithm.

Zusammenfassung

Diese Dissertation besteht aus zwei unabhängigen Teilen. Im ersten Teil entwickeln wir auf der Arbeit von Baker und Norine (*Advances in Mathematics*, 215(**2**): 766-788, 2007) aufbauend eine Riemann-Roch Theorie für Untergitter (sublattices) des Wurzelgitter (root lattice) A_n und untersuchen die Fragestellungen, die sich daraus ergeben. Unsere Theorie basiert auf der Untersuchung kritischer Punkte einer bestimmten simplizialen (simplicial) Metrik (distance function) auf Gitter und zeigt Verbindungen zwischen der Riemann-Roch Theorie und Voronoi-Diagrammen von Gittern unter einer gewissen simplizialen Metrik. Insbesondere liefern wir einen neuen geometrischen Beweis des Riemann-Roch Theorems für Graphen und generalisieren das Resultat für andere Untergitter von A_n. Des Weiteren verwenden wir den geometrischen Ansatz um das Problem der Berechnung des Rang (rank) eines Teilers (divisor) auf einem endlichen Multigraphen G und erhalten einen Algorithmus, der für eine fixe Anzahl von Knoten in Polynomialzeit, genauer in Zeit $2^{O(n \log n)}\text{poly}(\text{size}(G))$ mit n ist die Anzahl der Knoten in G, läuft. Von dieser Theorie ausgehend untersuchen wir einen Anzatz für das Graphautomorphismusproblem über eine Dimensionalitätsreduktion und erhalten ebenfalls einen Algorithmus für das verwandte Problem des Zählens von Automorphismen eines Graphen, der auf exponentiellen Summen basiert.

Im zweiten Teil der Dissertation entwickeln wir einen auf komplexwertigen Hashfunktionen basierenden Ansatz um in einem Streaming-Modell die Zyklen eines Graphen zu zählen. Unser Algorithmus basiert auf der Idee Instanzen von komplexwertigen Zufallsvariablen über dem gegebenen Stream zu berechnen und stellt eine drastische Verbesserung über den naiven Sampling-Algorithmus dar.

Contents

Introduction

This thesis consists of two independent parts. The first part is devoted to developing a Riemann-Roch theory on sublattices of the root lattice A_n and towards studying problems that are motivated by this theory. In the second part, we develop an algebraic approach to counting cycles in the datastreaming model. Chapters 1 and 2 are based on joint work with Omid Amini ([3]), Chapter 6 is based on joint work with Vikram Sharma ([61]) and Chapter 8 is based on joint work with Kurt Mehlhorn, Konstantinos Panagiotou and He Sun [60]. In this chapter, we discuss the motivation and context of our work and, then summarise the contributions of the thesis.

I. Riemann-Roch Theory for Sublattices of A_n

We start with a brief exposition on the history and motivation behind Riemann-Roch theory in discrete mathematics. There has been a surge of recent activity in promoting the viewpoint of a graph as the discrete analogue of a Riemann surface. In this context, the Laplacian matrix of a graph plays the key role of the Laplacian operator on a Riemann surface and there has been a lot of work on developing combinatorial analogues of results on Riemann surfaces. The article "Discrete Geometric Analysis" [79] of Sunada gives a broad survey of this perspective. Towards this direction, Baker and Norine [12] pioneered analogues of the Riemann-Roch theory in algebraic geometry and complex analysis in the context of graphs. Subsequently, the Rieman-Roch theorem on graphs played a key role in the development of a Riemann-Roch theory for tropical curves, see [63] and [39] for more details on the connections to tropical geometry. For a finite graph, the theory is best described in the language of chip-firing games.

A chip-firing game is a solitary game played on an undirected connected graph and is defined as follows: Each vertex of the graph is assigned an integer (positive or negative), refered to as "the number of chips" and this assignment of chips is called the initial configuration. At each move of the game, an arbitrary vertex v is allowed to either lend or borrow one chip along each incident edge resulting in a new configuration. We define two configurations C_1 and C_2 to be *linearly equivalent* if C_1 can be reached from C_2 by a sequence of chip firings. The Laplacian matrix Q of the graph naturally comes into the picture as follows:

Lemma 1. *Configurations C_1 and C_2 are linearly equivalent if and only if $C_1 - C_2$ can be expressed as $Q \cdot w$ for some vector w with integer coordinates.*

We can ask some natural questions on such a game:

1. Is a given configuration linearly equivalent to an *effective* configuration, i.e., a configuration where each vertex has a non-negative number of chips?

2. More generally, given a configuration, what is the minimum number of chips that must be removed from the system so that the resulting configuration is not linearly equivalent to an effective configuration?

The Riemann-Roch theorem of Baker and Norine provides insights into answering these questions. We need a couple of definitions before we can state the theorem. For a configuration $\mathcal{C}$, one less than the minimum number of chips that must be removed from the configuration $\mathcal{C}$ so that the resulting configuration is not linearly equivalent to an effective configuration is called the rank of the configuration $\mathcal{C}$, denoted by $r(\mathcal{C})$. For a configuration $\mathcal{C}$, the total number of chips, i.e., the sum of chips over all the vertices of the graph, is called the degree of the configuration $\mathcal{C}$, denoted by $\deg(\mathcal{C})$.

Theorem 1. (**Riemann-Roch theorem for graphs** [12]) *For any undirected connected graph G, there exists a special configuration $\mathcal{K}$ called the canonical divisor such that for any configuration $\mathcal{C}$ we have*

$$r(\mathcal{C}) - r(\mathcal{K} - \mathcal{C}) = \deg(\mathcal{C}) - g + 1. \tag{1}$$

Here g is the genus of G (also known as the cyclotomic number of G) and is equal to $m - n + 1$ where m is the number of edges and n is the number of vertices of G.

The proof of Theorem 1 by Baker and Norine is combinatorial and the main component of the proof is the existence and uniqueness of certain *normal* forms called reduced divisors.

Our work started with the observation that the set-up can be generalised to an arbitrary full-rank sublattice of the root lattice A_n defined as $A_n = \{(x_1, \ldots, x_{n+1}) | \sum_{i=1}^{n+1} x_i = 0, x_i \in \mathbb{Z}\}$ and that the Riemann-Roch theorem for graphs can be rephrased in terms of the lattice generated by the rows of the Laplacian matrix of a graph. We call this lattice the Laplacian lattice of a graph. This led to the question of whether a Riemann-Roch type theorem is true for an arbitrary full-rank sublattice of A_n. A computer check showed that this is not true in general. Hence, the problem of characterizing the geometric properties of a sublattice of A_n for which the Riemann-Roch theorem holds arises. In fact, one of our main results is a characterization of sublattice of A_n that satisfy a Riemann-Roch formula. We now briefly set up the reformulation of the Riemann-Roch theory for graphs in terms of the Laplacian lattice and then state our main results.

Reformulation in Terms of the Laplacian Lattice. Recall that a lattice is a discrete subgroup of the group $(\mathbb{R}^n, +)$ for some integer n (for instance, the lattice $\mathbb{Z}^n \subset \mathbb{R}^n$), and the rank of a lattice is its rank considered as a free Abelian group. We call a sublattice of $\mathbb{Z}^n$ an integral lattice.

Let $G = (V, E)$ be a given undirected connected (multi-)graph and $V = \{v_0, \dots, v_n\}$. The *Laplacian* of G is the matrix $Q = D - A$, where D is the diagonal matrix whose $(i,i)-$th entry is the degree of v_i, and A is the *adjacency matrix* of G whose $(i,j)-$th entry is the number of edges between v_i and v_j. Some basic facts about Q are that it is symmetric and has rank n for a connected graph G, and that the kernel of Q is spanned by the vector whose entries are all equal to 1, cf. [16].

The *Laplacian lattice* L_G of G is defined as the image of $\mathbb{Z}^{n+1}$ under the linear map defined by Q, i.e., $L_G := Q(\mathbb{Z}^{n+1})$, c.f., [10]. Since G is a connected graph, L_G is a sublattice of the root lattice A_n of full-rank equal to n, where $A_n \subset \mathbb{R}^{n+1}$ is the lattice defined as follows[1]:

$$A_n := \Big\{ x = (x_0, \dots, x_n) \in \mathbb{Z}^{n+1} \mid \sum x_i = 0 \Big\}.$$

Note that A_n is a discrete sub-group in the hyperplane

$$H_0 = \Big\{ x = (x_0, \dots, x_n) \in \mathbb{R}^{n+1} | \sum x_i = 0 \Big\}$$

of $\mathbb{R}^{n+1}$ and has rank n.

Once we fix a labelling of the vertices of G, it is straightforward to associate a point $D_{\mathcal{C}}$ in $\mathbb{Z}^{n+1}$ to each configuration $\mathcal{C}$ as follows: $D_{\mathcal{C}}$ is the vector with coordinates equal to the number of chips given to the vertices of G. For a sequence of chip-firings on $\mathcal{C}$ resulting in another configuration $\mathcal{C}'$, then there exists a vector $v \in L_G$ such that $D_{\mathcal{C}'} = D_{\mathcal{C}} + v$. Conversely, if $D_{\mathcal{C}'} = D_{\mathcal{C}} + v$ for a vector $v \in L_G$, then there is a sequence of chip-firings transforming $\mathcal{C}$ to $\mathcal{C}'$. Using this equivalence, we transform the chip-firing game and the statement of the Riemann-Roch theorem to a statement about $\mathbb{Z}^{n+1}$ and the Laplacian lattice $L_G \subseteq A_n$.

Remark 1. (**The Laplacian matrix of a graph**) The Laplacian matrix of graph and its spectral theory have been well studied. The Laplacian matrix captures information about the geometry and combinatorics of the graph G, for example, it provides bounds on the expansion of G (we refer to the survey [48]) or on the quasi-randomness properties of the graph, see [28]. The famous Matrix Tree Theorem states that *the cardinality of the (finite) Picard group* $\mathrm{Pic}(G) := A_n/L_G$ *is the number of spanning trees of* G. We refer to classic books on algebraic graph theory such as [41] and [16] for more details.

Remark 2. (**Lattices associated with graphs**) A substantial body of work is devoted to the study of lattices constructed from graphs, the most well studied ones being the lattice of integral cuts and the lattice of integral flows. This line of investigation was pioneered by the work of Bacher, Harpe and Nagnibeda [10] where they provide a combinatorial interpretation of various parameters of the lattice of integral flows and the lattice of integral cuts in the Euclidean distance function, for example they show that the square norm of the shortest vector of the lattice of integral flows is equal to the girth of the graph.

[1]Root refers here to root systems in the classification theory of simple Lie algebras [21]

Linear Systems of Integral Points and the Rank Function. Let L be a sublattice of A_n of full-rank (for example, $L = L_G$ for a finite graph G). Define an equivalence relation $\sim$ on the set of points of $\mathbb{Z}^{n+1}$ as follows: $D \sim D'$ if and only if $D - D' \in L$. This equivalence relation is referred to as *linear equivalence* and the equivalence classes are denoted by $\mathbb{Z}^{n+1}/L$. We say that a point E in $\mathbb{Z}^{n+1}$ is *effective* or *non-negative*, if all the coordinates are non-negative. For a point $D \in \mathbb{Z}^{n+1}$, the linear system associated to D is the set $|D|$ of all effective points linearly equivalent to D:

$$|D| = \Big\{ E \in \mathbb{Z}^{n+1} \,:\, E \geq 0,\;\; E \sim D \Big\}.$$

The *rank* of an integral point $D \in \mathbb{Z}^{n+1}$, denoted by $r(D)$, is defined by setting $r(D) = -1$, if $|D| = \emptyset$, and then declaring that for each integer $s \geq 0$, $r(D) \geq s$ if and only if $|D - E| \neq \emptyset$ for all effective integral points E of degree s. Observe that $r(D)$ is well-defined and only depends on the linear equivalence class of D. Note that $r(D)$ can be defined as follows:

$$r(D) = \min \Big\{ \deg(E) \mid |D - E| = \emptyset,\; E \geq 0 \Big\} - 1.$$

Obviously, $\deg(D)$ is a trivial upper bound for $r(D)$.

Remark 3. A natural question that arises with the set-up of the Riemann-Roch machinery on a finite graph is its connection to the classical Riemann-Roch machinery on a Riemann surface, more generally on an algebraic curve and this question is addressed in Baker's Specialization Lemma [11]: the rank of a divisor on an algebraic curve is upper bounded by its "specialisation" to the rank of a divisor on a finite graph, see [11] for a more precise statement.

Our Contributions

Extension of the Riemann-Roch Theorem to Sublattices of A_n. In Chapter 1, we provide a characterization of the sublattices of A_n which admit a Riemann-Roch theorem with respect to the rank-function defined above. Furthemore, our approach provides a geometric proof of the Riemann-Roch theorem of Baker and Norine (Theorem 1).

We show that Riemann-Roch theory associated to a full rank sublattice L of A_n is related to the study of the Voronoi diagram of the lattice L in the hyperplane H_0 under a certain simplicial distance function which we refer to in this chapter as "the simplicial distance function"[2]. The whole theory is then captured by the corresponding critical points of this simplicial distance function.

[2]Let S be a full dimensional simplex in $\mathbb{R}^n$ that contains the origin, that we call the center of S, in its interior. For two points $P_1,\ P_2 \in \mathbb{R}^n$, the simplicial distance function induced by S is the smallest factor by which we need to scale the translated copy of S centered at P_1 so that it contains P_2. The Euclidean distance can be obtained by replacing the simplex by a sphere centered at the origin. Given a lattice, the simplicial distance function of a point with respect to a lattice is the minimum simplicial distance of the point to any lattice point.

We associate two geometric invariants to each such sublattice of A_n, the *min*- and the *max-genus*, denoted respectively by g_{min} and g_{max}. Two main characteristic properties for a given sublattice of A_n are then defined. The first one is what we call *Reflection Invariance* (Definition 1.5.1 of Chapter 1), and one of our results here is a weak Riemann-Roch theorem for reflection-invariant sublattices of A_n of full-rank n.

Theorem 2. (Weak Riemann-Roch)(Theorem 1.5.2) *Let L be a reflection invariant sublattice of A_n of rank n. There exists a point $K \in \mathbb{Z}^{n+1}$, called* canonical point*, such that for every point $D \in \mathbb{Z}^{n+1}$, we have*

$$3g_{min} - 2g_{max} - 1 \;\leq\; r(K-D) - r(D) + \deg(D) \;\leq\; g_{max} - 1\,.$$

The second characteristic property is called *Uniformity* and simply means $g_{min} = g_{max}$. It is straightforward to derive a Riemann-Roch theorem for uniform reflection-invariant sublattices of A_n of rank n from Theorem 2 above.

Theorem 3. (Riemann-Roch) *Let L be a uniform reflection invariant sublattice of A_n. Then there exists a point $K \in \mathbb{Z}^{n+1}$, called* canonical*, such that for every point $D \in \mathbb{Z}^{n+1}$, we have*

$$r(D) - r(K-D) = \deg(D) - g + 1,$$

where $g = g_{min} = g_{max}$.

We then show, in Chapter 2, that Laplacian lattices of undirected connected graphs are uniform and reflection invariant, obtaining a geometric proof of the Riemann-Roch theorem for graphs. As a consequence of our results, we provide an explicit description of the Voronoi diagram of lattices generated by Laplacians of connected graphs. In the case of Laplacian lattices of connected regular digraphs, we also provide a slightly stronger statement than the weak Riemann-Roch Theorem (Theorem 2) above. The results lead us into two natural directions: (i) Obtaining a complete understanding of the structure of the Laplacian lattice under the simplicial distance function and (ii) Algorithmic aspects.

The Laplacian lattice of a Graph under the Simplicial Distance Function. In Chapter 3, we describe important geometric invariants of the Laplacian lattice under the simplicial distance function, namely the minimal vectors, the covering and packing radius, the local maxima of the simplicial distance function and the Delaunay triangulations. It turns out that most of these invariants have combinatorial interpretations in terms of parameters of the underlying graph.

In Chapter 4 we study some applications of the combinatorial interpretation of the geometric invariants of the Laplacian lattice. A natural question that arises with the construction of the Laplacian lattice is whether it determines the underlying graph completely up to isomorphism. More precisely, are Laplacian lattices L_G and L_H congruent if and only if G is isomorphic to H. The answer to this question is "no" since, every tree on $n+1$ vertices has A_n as its Laplacian lattice. Nevertheless, for a graph G we associate a convex polyope $H_{Del_G}(O)$, which is in fact the Delaunay polytope of L_G under the simplicial distance function and show that $H_{Del_G}(O)$ characterizes the graph completely up to isomorphism. More precisely,

Theorem 4. (Theorem 4.1.20) *Let G_1 and G_2 be undirected connected graphs. The polytopes $H_{Del_{G_1}}(O)$ and $H_{Del_{G_2}}(O)$ are congruent if and only if G_1 and G_2 are isomorphic.*

As we prove Theorem 4.1.20, we undertake a detailed study of the combinatorial structure of the polytope $H_{Del_G}(O)$, determining its vertices, edges and facets. Another related question is to count graphs with a given lattice as their Laplacian lattice. In fact by the Matrix-Tree theorem we know that for the root lattice A_n, this number is equal to $(n+1)^{n-1}$, the number of spanning trees of the complete graph K_{n+1}. We show the following generalisation of this observation:

Theorem 5. (Theorem 4.2.2) *For a Laplacian lattice L_G, let $N_{\mathrm{Gr}}(L_G)$ be the number of graphs with L_G as their Laplacian lattice and let $N_{\mathrm{Del}}(L_G, \triangle)$ be the number of different Delaunay triangulations of L_G under the simplicial distance function, we have $N_{\mathrm{Gr}}(L_G) \leq N_{\mathrm{Del}}(L_G, \triangle)$.*

Furthermore, we use the combinatorial interpretation of the geometric invariants of the Laplacian lattice to relate the connectivity properties of the graph to the covering and packing density of the corresponding Laplacian lattice in the simplicial distance function. In particular, we show that in the space of Laplacian lattices of undirected connected graphs, Laplacian lattices of graphs with high-connectivity such as Ramanujan graphs possess good packing and covering density. More precisely, we show the following:

Theorem 6. (Theorem 4.3.6) *Let G be a d-regular Ramanujan graph, then the covering density of L_G under the simplicial distance function is at most $\frac{d}{4(d-2\sqrt{d-1})}$.*

Theorem 7. (Theorem 4.3.8) *Let G be a d-regular Ramanujan graph, then the packing density of L_G under the simplicial distance function is at least $\frac{(d-2\sqrt{d-1})}{2(n+1)(d+2\sqrt{d-1})}$.*

The above result is also motivated by the fact that explicit constructions of lattices with optimal packing and covering densities are known only in some small dimensions [31].

Algorithmic Aspects. A central quantity in the Riemann-Roch theorem is the rank of a divisor and the efficient computation of rank is a natural question. The problem of deciding if the rank of a divisor on a finite multigraph is non-negative has been considered by several authors independently and the problem can be solved in polynomial time in the size of the input, see [32], [80] and [13]. On the other hand, the problem of computing the rank of a divisor seems to be harder in terms of computational complexity. Hladkỳ, D. Král and Norine [46] construct an algorithm, i.e., a procedure that terminates in a finite number of steps, to compute the rank of a divisor on a tropical curve and the algorithm also computes the rank of a divisor on a finite multigraph. However, the algorithm does not run in polynomial time in the number of bits needed to encode the

Laplacian matrix of the multigraph even when the number of vertices are fixed since the algorithm involves iterating over all the spanning trees in the graph (see proof of Theorem 23 in [46]) and the number of spanning trees in not polynomially bounded in the size of the multigraph even if the number of vertices are fixed. We obtain an algorithm (Chapter 5) for computing the rank of a divisor on a finite multigraph that runs in polynomial time for a fixed number of vertices. More precisely,

Theorem 8. (Theorem 5.1.17) *There is an algorithm that computes the rank of a divisor on a finite multigraph G on n vertices with running time $2^{O(n \log n)} poly(size(G))$ and hence, runs in time polynomial in the size of the input for a fixed number of vertices.*

The key ingredients of the algorithm include a new geometric interpretation of rank (Theorem 5.1.13) combined with Kannan's algorithm for integer programming from the geometry of numbers [51]. On the other hand, we show that computing the rank of a configuration on a general sublattice of A_n i.e., the corresponding decision problem is NP-hard.

The correspondence between lattices and graphs raises the question of whether we can use geometric methods to better understand computational problems on graphs; one such example is the Graph Isomorphism problem. The Graph Isomorphism problem, i.e. the computational problem of testing if two graphs are isomorphic, is a standard example of a problem in the complexity class NP that is neither known to have a polynomial time algorithm nor is known to be NP-hard. We refer to [54] for a detailed survey on the complexity theoretic aspects of the problem. Another problem closely related to the Graph Isomorphism problem is the Graph Automorphism problem, the computational problem of testing if a graph has a non-trivial automorphism. It is another example of a problem in NP that is neither known to have a polynomial time algorithm nor known to be NP-hard. In Chapter 6, we focus on the graph automorphism problem.

We start with the result that we obtained while studying Laplacian lattices namely that the polytope of the Laplacian lattice with respect to the simplicial distance function characterizes the graph completely up to isomorphism. We further observe that this result can be "simplified" in the following sense: the Laplacian simplex, i.e., the convex hull of the rows of the Laplacian matrix, also characterizes the graph completely up to isomorphism. This observation raises the question of whether we can apply high-dimensional geometry techniques such as dimensionality-reduction to solve the graph automorphism problem efficiently. In particular, can we project the vertices of the simplex onto a low-dimensional space and then use a congruence testing algorithm to test the "approximate congruence" of the projected point sets? A priori, even the existence of such low-dimensional spaces is not clear. We show (Theorem 6.1.9) the existence of low-dimensional subspaces that preserve graph automorphisms and characterize them as certain low-dimensional invariant subspaces of automorphisms that we call "cross-invariant subspaces". The question of efficient construction of these spaces, however, is still open.

In the second part of Chapter 6, we provide an exponential sum formula to count the number of automorphisms of a graph and study its complexity. Let us now briefly outline the construction of the exponential sum formula. Given a graph G, we start by constructing a function on the set of permutations σ on the vertex set of G with the following property: the function vanishes for a permutation if and only if σ is an automorphism of G. The desired function is:

$$f(\sigma) = \frac{1}{2}\sum_{i=1}^{n+1}\sum_{j=1}^{n+1}(a_{i,j} - a_{\sigma(i),\sigma(j)})^2; \tag{2}$$

where $a_{i,j}$ is an entry in the vertex to vertex adjacency matrix of G. It follows that $f(\sigma)$ has the following property:

Lemma 2. *If σ is an automorphism of G, then the function $f(\sigma) = 0$; otherwise $f(\sigma)$ counts the number of edges violated by the permutation σ, i.e., the number of edges that are mapped to non-edges and vice-versa.*

Thus, we can interpret $f(\sigma)$ as an indicator function for being an automorphism over the set of all permutations. Suppose $f(\sigma)$ was equal to some c for all non-automorphisms then the quantity $\sum_{\sigma\in S_n} f(\sigma)/c$ will give us the number of non-automorphisms, and the number of automorphisms can also be computed from this information. However, this assumption may not be true in general. To salvage this approach, we use the following property of exponential sums: For an integer m and a prime p we know that

$$\sum_{k=0}^{p-1}\exp(2\pi ikm/p) = \begin{cases} p & \text{if } p \mid m, \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

Using this property we have the following formula to count the number of automorphisms.

Theorem 9. (Theorem 6.2.5) *For a sufficiently large prime p, the number of automorphisms N_A, is equal to*

$$\frac{1}{p}\sum_{\sigma\in S_n}\sum_{k=0}^{p-1}\exp(2\pi ikf(\sigma)/p). \tag{4}$$

In particular, we can choose p to be larger than $\max f(\sigma)$ over all $\sigma \in S_n$.

We then study the computation of the Formula (4) and show that the lower order terms of its Taylor series expansion can be computed efficiently. As a consequence, we obtain the following algorithmic result:

Theorem 10. (Theorem 6.2.7) *For any graph G and a fixed prime p, there is an algorithm that counts, modulo p, the number of permutations that violate a multiple of p edges in G, and the running time of the algorithm is polynomial in the size of the input.*

This is perhaps the best that can be done in polynomial time, since it is known that counting the number of automorphisms modulo two is at least as hard as Graph Automorphism, and also counting the exact number of automorphisms is at least as hard as Graph Isomorphism [6].

II. Counting Cycles in the Data Streaming Model

We now introduce the second part of the thesis where we develop an approach to counting subgraphs based on complex-valued hash functions, in particular to counting cycles of a graph presented as a data stream. Counting the number of occurrences of a graph H in a graph G has wide applications in uncovering important structural characteristics of the underlying network G, for example in revealing information of the most frequent patterns. We are particularly interested in the situation where G is very large. It is then natural to assume that G is given as a data stream, i.e., the edges of the graph G arrive consecutively and the algorithm uses only limited space to return an approximate value. Exact counting is not an option for massive input graphs. Indeed counting triangles exactly already requires us to store the entire graph.

We now briefly formalise our model of computation. Let $\mathcal{S} = s_1, s_2, \cdots, s_N$ be a stream that represents a graph $G = (V, E)$, where N is the length of the stream and each item s_i is associated with an edge in G. Typical models [66] in this topic include *the Cash Register Model* and *the Turnstile Model*. In the cash register model, each item s_i expresses one edge in G, and in the turnstile model each item s_i is represented by (e_i, s_i) where e_i is an edge of G and $s_i \in \{+, -\}$ indicates that e_i is inserted to or deleted from G. As a generalization of the cash register model, the turnstile model supports dynamic insertions and deletions of the edges.

In a *distributed* setting the stream $\mathcal{S}$ is partitioned into sub-streams $S_1, \ldots, S_t$ and each S_i is fed to a different processor. At the end of the computation, the processors collectively estimate the number of occurrences of H with a small amount of communication.

Counting subgraphs in a data stream was first considered in a seminal paper by Bar-Yossef, Kumar, and Sivakumar [81]. There, the triangle counting problem was reduced to the problem of computing frequency moments. After that, several algorithms for counting triangles have been proposed [49].

Jowhari and Ghodsi presented three algorithms in [49], one of which is applicable in the turnstile model. Moreover, the problem of counting subgraphs different from triangles has also been investigated in the literature. Bordino, Donato, Gionis, and Leonardi [20] extended the technique of counting triangles [22] to all subgraphs on three and four vertices. Buriol, Fahling, Leonardi and Sohler [56] presented a streaming algorithm for counting $K_{3,3}$, the complete bipartite graph with three vertices in each part. However, except the one presented in [49], most algorithms are based on sampling techniques and do not apply to the turnstile model.

Our Contributions

We present a general framework for counting cycles of arbitrary size in a massive graph. Our algorithm runs in the turnstile model and the distributed setting, and for any constants $0 < \varepsilon, \delta < 1$, our algorithm achieves an (ε, δ)-approximation, i.e., the output Z of the algorithm and the exact value $Z^* = \#C_k$, the number of occurrences of C_k, satisfy $\Pr[|Z - Z^*| > \varepsilon \cdot Z^*] < \delta$. We also provide an unbiased estimator for general d-regular graphs. This considerably extends the class of graphs that can be counted in the data streaming model and answers partially an open problem proposed by many references, see for example the extensive survey by Muthukrishnan [66] and the 11th open question in the 2006 IITK Workshop on Algorithms for Data Streams [69]. Besides that, we initiate the study of complex-valued hash functions in counting subgraphs.

Remark 4. Complex-valued estimators have been successfully applied in other contexts such as approximating the permanent, see [53]. In the data streaming setting, Ganguly [38] used a complex-valued sketch to estimate frequency moments.

Our main result is as follows:

Theorem 11. (Theorem 8.3.5) *Let G be a graph with n vertices and m edges. For any k, there is an algorithm using S bits of space to (ε, δ)-approximate the number of occurrences of C_k in G provided that $S = \Omega\big(\frac{1}{\varepsilon^2} \cdot \frac{m^k}{(\#C_k)^2} \cdot \log n \cdot \log \frac{1}{\delta}\big)$. The algorithm works in the turnstile model.*

A naïve approach for counting the number of occurrences of a k-cycle would either sample independently k vertices (if possible) or k edges from the stream. Since the probability of k vertices (or k edges) forming a cycle is $\#C_k/n^k$ (or $\#C_k/m^k$), this approach needs space $\Omega\left(\frac{n^k \log n}{\#C_k}\right)$ and $\Omega\left(\frac{m^k \log n}{\#C_k}\right)$, respectively. Thus, our algorithm improves upon these two approaches, especially for sparse graphs with many k-cycles, and has the additional benefit that it is applicable in the turnstile model and the distributed setting. Moreover, note that our bound is essentially tight, as there are graphs where the space complexity of the algorithm is $O(\log n)$: consider for example the "extremal graph" with a clique on $\Theta(\sqrt{m})$ vertices, where all other vertices are isolated. Moreover, as a corollary of Theorem 11, when the number of occurrences of C_k is $\Omega\left(m^{k/2-\alpha}\right)$ for $0 \leq \alpha < 1/2$, our algorithm with sub-linear space $O\left(\frac{1}{\varepsilon^2} \cdot m^{2\alpha} \cdot \log \frac{1}{\delta}\right)$ suffices to give a good approximation.

Part I

Riemann-Roch Theory for Sublattices of A_n

Chapter 1

Riemann-Roch for Sublattices of A_n

In this chapter, we establish a Riemann-Roch theorem for sublattices of the root lattice A_n. Let us start by briefly discussing the notation that we will frequently use in this chapter as well as in the upcoming chapters.

1.1 Basic Notations

A point of $\mathbb{R}^{n+1}$ with integer coordinates is called an *integral point*. By a lattice L, we mean a discrete subgroup of H_0 of maximum rank. Recall that H_0 is the set of all points of $\mathbb{R}^{n+1}$ such that the sum of their coordinates is zero. The elements of L are called *lattice points*. The positive cone in $\mathbb{R}^{n+1}$ consists of all the points with non-negative coordinates. We can define a partial order in $\mathbb{R}^{n+1}$ as follows: $a \leq b$ if and only if $b - a$ is in the positive cone, i.e., if each coordinate of $b - a$ is non-negative. In this case, we say that b *dominates* a. Also we write $a < b$ if all the coordinates of $b - a$ are strictly positive.

For a point $v = (v_0, \ldots, v_n) \in \mathbb{R}^{n+1}$, we denote by v^- and v^+ the negative and positive parts of v respectively. For a point $p = (p_0, \ldots, p_n) \in \mathbb{R}^{n+1}$, we define the degree of p as $\deg(p) = \sum_{i=0}^{n} p_i$. For each k, by H_k we denote the hyperplane consisting of points of degree k, i.e., $H_k = \{x \in \mathbb{R}^{n+1} \mid \deg(x) = k\}$. By π_k, we denote the projection from $\mathbb{R}^{n+1}$ onto H_k along $\vec{\mathbf{1}} = (1, \ldots, 1)$. In particular, π_0 is the projection onto H_0. Finally for an integral point $D \in \mathbb{Z}^{n+1}$, by $N(D)$ we denote the set of all neighbours of D in $\mathbb{Z}^{n+1}$, which consists of all the points of $\mathbb{Z}^{n+1}$ which have distance at most one to D in ℓ_1 norm.

In the following, to simplify the presentation, we will use the convention of tropical arithmetic, briefly recalled below. The tropical semiring $(\mathbb{R}, \oplus, \otimes)$ is defined as follows: As a set this is just the real numbers $\mathbb{R}$. However, one redefines the basic arithmetic operations of addition and multiplication of real numbers as follows: $x \oplus y := \min(x, y)$ and $x \otimes y := x + y$. In words, the tropical sum of two numbers is their minimum, and the tropical product of two numbers is their sum. Similarly we denote the maximum of two real numbers x, y by $x \oplus_{\max} y$. We extend the tropical sum, tropical product and the maximum to vectors by doing the operations coordinate-wise.

1.2 Preliminaries

All through this section L will denote a full rank integral sublattice of H_0 i.e., a sublattice of A_n.

1.2.1 Sigma-Region of a Given Sublattice L of A_n

Every point D in $\mathbb{Z}^{n+1}$ defines two "orthogonal" cones in $\mathbb{R}^{n+1}$, denoted by H_D^- and H_D^+, as follows: H_D^- is the set of all points in $\mathbb{R}^{n+1}$ which are dominated by D. In other words

$$H_D^- = \{\, D' \mid D' \in \mathbb{R}^{n+1}, D - D' \geq 0 \,\}.$$

Similarly H_D^+ is the set of points in $\mathbb{R}^{n+1}$ that dominate D. In other words,

$$H_D^+ = \{\, D' \mid D' \in \mathbb{R}^{n+1}, D' - D \geq 0 \,\}.$$

For a cone $\mathcal{C}$ in $\mathbb{R}^{n+1}$, we denote by $\mathcal{C}(\mathbb{Z})$ and $\mathcal{C}(\mathbb{Q})$, the set of integral and rational points of the cone respectively. When there is no risk of confusion, we sometimes drop $(\mathbb{Z})$ (resp. $(\mathbb{Q})$) and only refer to $\mathcal{C}$ as the set of integral points (resp. rational points) of the cone $\mathcal{C}$. The *Sigma-Region* of the lattice L is, roughly speaking, the set of integral points of $\mathbb{Z}^{n+1}$ that are not contained in the cone H_p^- for any point $p \in L$. More precisely:

Definition 1.2.1. *The Sigma-Region of L, denoted by $\Sigma(L)$, is defined as follows:*

$$\Sigma(L) = \{\, D \mid D \in \mathbb{Z}^{n+1} \;\&\; \forall\, p \in L,\; D \not\leq p \,\} = \mathbb{Z}^{n+1} \setminus \bigcup_{p \in L} H_p^-. \tag{1.1}$$

The following lemma shows the relation between the Sigma-Region and the rank of an integral point as defined in the previous section.

Lemma 1.2.2.

(i) For a point D in $\mathbb{Z}^{n+1}$, $r(D) = -1$ if and only if $-D$ is a point in $\Sigma(L)$.

(ii) More generally, $r(D) + 1$ is the distance of $-D$ to $\Sigma(L)$ in the ℓ_1 norm, i.e.,

$$r(D) \;= dist_{\ell_1}(-D, \Sigma(L)) - 1 := \inf\{||p + D||_{\ell_1} \mid p \in \Sigma(L)\} - 1,$$

where $||x||_{\ell_1} = \sum_{i=0}^{n} |x_i|$ for every point $x = (x_0, x_1, \ldots, x_n) \in \mathbb{R}^{n+1}$.

Before presenting the proof of Lemma 1.2.2, we need the following simple observation.

Observation 1. *$\forall\, D_1, D_2 \in \mathbb{Z}^{n+1}$, we have $D_1 \in \Sigma(L) - D_2$ if and only if $D_2 \in \Sigma(L) - D_1$.*

We shall usually use this observation without sometimes mentioning it explicitly.

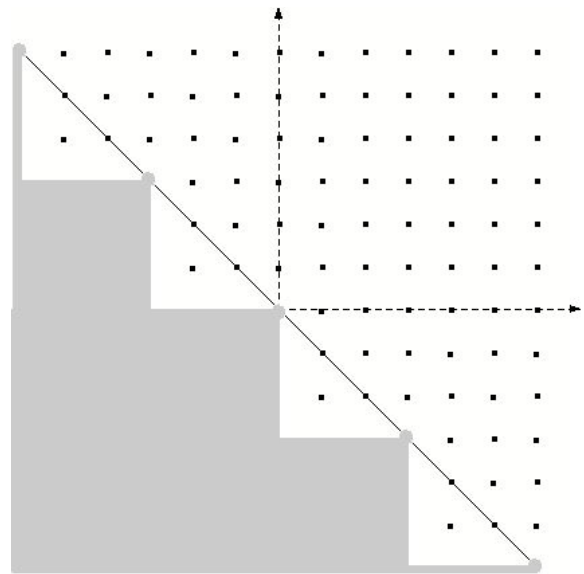

Figure 1.1: A finite portion of the Sigma-Region of a sublattice of A_1. All the black points belong to the Sigma-Region. The integral points in the grey part are out of the Sigma-Region.

Proof of Lemma 1.2.2.

(i) Recall that $r(D) = -1$ means that $|D| = \emptyset$. This in turn means that $D \not\geq p$ for any p in L, or equivalently $-D \not\leq q$ for any point q in L (because $L = -L$). We infer that $-D$ is a point of $\Sigma(L)$. Conversely, if $-D$ belongs to $\Sigma(L)$, then $-D \not\leq q$ for any point q in L, or equivalently $D \not\geq p$ for any p in L (because $L = -L$). This implies that $|D| = \emptyset$ and hence $r(D) = -1$.

(ii) Let p^* be a point in $\Sigma(L)$ which has minimum ℓ_1 distance from $-D$, and define $v^* = p^* + D$. Write $v^* = v^{*,+} + v^{*,-}$, where $v^{*,+}$ and $v^{*,-}$ are respectively the positive and the negative parts of v^*. We first claim that v^* is an effective integral point, i.e., $v^{*,-} = 0$. For the sake of a contradiction, let us assume the contrary, i.e., assume that $||v^{*,-}||_{\ell_1} > 0$. Since $-D + v^{*,+} + v^{*,-} = -D + v^* = p^*$ is contained in $\Sigma(L)$, and because $v^{*,-} \leq 0$, the point $p^{*,+} = -D + v^{*,+}$ has to be in $\Sigma(L)$. Also $||v^{*,+}||_{\ell_1} < ||v^*||_{\ell_1}$ (because $||v^*||_{\ell_1} = ||v^{*,+}||_{\ell_1} + ||v^{*,-}||_{\ell_1}$ and $||v^{*,-}||_{\ell_1} > 0$). We obtain $||D + p^{*,+}||_{\ell_1} = ||v^{*,+}||_{\ell_1} < ||D + p^*||_{\ell_1}$, which is a contradiction by the choice of p^*. Therefore, we have

$$\begin{aligned}
r(D) &= \min\{\, \deg(v) \mid |D-v| = \emptyset,\ v \geq 0 \,\} - 1 \\
&= \min\{\, \deg(v) \mid v - D \in \Sigma(L),\ v \geq 0 \,\} - 1 \quad \text{(By the first part of Lemma 1.2.2)} \\
&= \min\{\, ||v||_{\ell_1} \mid v - D \in \Sigma(L),\ v \geq 0 \,\} - 1 \\
&= \min\{\, ||D+p||_{\ell_1} - 1 \mid p \in \Sigma(L) \text{ and } D + p \geq 0 \,\} \\
&= dist_{\ell_1}(-D, \Sigma(L)) - 1 \qquad \text{(By the above arguments).}
\end{aligned}$$

□

Lemma 1.2.2 shows the importance of understanding the geometry of the Sigma-Region for the study of the rank function. This will be our aim in the rest of this section and in Section 1.4. But we need to introduce another definition before we proceed. Apparently, it is easier to work with a "continuous" and "closed" version of the Sigma-Region.

Definition 1.2.3. $\Sigma^{\mathbb{R}}(L)$ *is the set of points in* $\mathbb{R}^n$ *that are not dominated by any point in* L.

$$\begin{aligned}
\Sigma^{\mathbb{R}}(L) &= \Big\{\, p \mid p \in \mathbb{R}^{n+1} \ and\, p \not\leq q,\ \forall q \in L \,\Big\} \\
&= \mathbb{R}^{n+1} \setminus \bigcup_{p \in L} H_p^-.
\end{aligned}$$

By $\Sigma^c(L)$ *we denote the topological closure of* $\Sigma^{\mathbb{R}}(L)$ *in* $\mathbb{R}^{n+1}$.

Remark 1.2.4. One advantage of this definition is that it can be used to define the same Riemann-Roch machinery for any full rank sublattice of H_0. Indeed for such a sublattice L, it is quite straightforward to associate a real-valued rank function to any point of $\mathbb{R}^{n+1}$. The main theorems of this thesis can be proved in this more general setting. As all the examples of interest for us are integral lattices (lattices whose elements all have integer coordinates), we have restricted the presentation to sublattices of A_n.

1.2.2 Extremal Points of the Sigma-Region

We say that a point $p \in \Sigma(L)$ is an *extremal* point if it is a local minimum of the degree function. In other words

Definition 1.2.5. *The set of extremal points of* L *denoted by* $\mathrm{Ext}(L)$ *is defined as follows:*

$$\mathrm{Ext}(L) := \{\nu \in \Sigma(L) \mid \deg(\nu) \leq \deg(q)\ \forall\ q \in N(\nu) \cap \Sigma(L)\}).$$

Recall that for every point $D \in \mathbb{Z}^{n+1}$, $N(D)$ *is the set of neighbours of* D *in* $\mathbb{Z}^{n+1}$, *which consists of all the points of* $\mathbb{Z}^{n+1}$ *which have distance at most one to* D *in* ℓ_1 *norm.*

We also define extremal points of $\Sigma^c(L)$ as the set of points that are local minimum of the degree function and denote it by $\mathrm{Ext}^c(L)$. Local minimum here is understood with respect to the topology of $\mathbb{R}^{n+1}$: x is a local minimum if and only if there exists an open ball B containing x such that x is the point of minimum degree in $B \cap \Sigma^c(L)$. The following theorem describes the Sigma-Region of L in terms of its extremal points.

Theorem 1.2.6. *Let L be a sublattice of A_n with full rank. Every point of the Sigma-Region dominates an extremal point. More precisely, $\Sigma(L) = \cup_{\nu \in \mathrm{Ext}(L)} H_\nu^+(\mathbb{Z})$. Recall that $H_\nu^+(\mathbb{Z})$ is the set of integral points of the cone H_v^+.*

Indeed, we first prove the following continuous version of Theorem 1.2.6.

Theorem 1.2.7. *For any (integral) sublattice L of H_0, we have $\Sigma^c(L) = \cup_{\nu \in \mathrm{Ext}^c(L)} H_\nu^+$.*

And Theorem 1.2.6 is derived as a consequence of Theorem 1.2.7. The proof of these two theorems are presented in Section 1.3. The proof shows that every extremal point of $\Sigma^c(L)$ is an integral point and $\Sigma(L) = \Sigma^c_{\mathbb{Z}}(L) + (1, \dots, 1)$, where $\Sigma^c_{\mathbb{Z}}(L)$ denotes the set of integral points of $\Sigma^c(L)$. We refer to Section 1.3 for more details.

Proposition 1.2.8. *We have $\Sigma(L) = \Sigma^c_{\mathbb{Z}}(L) + (1, \dots, 1)$ and $\mathrm{Ext}(L) = \mathrm{Ext}^c(L) + (1, \dots, 1)$. In particular, $\pi_0(\mathrm{Ext}^c(L)) = \pi_0(\mathrm{Ext}(L))$.*

The important point about Theorem 1.2.6 is that one can use it to express $r(D)$ in terms of the extremal points of $\Sigma(L)$. For an integral point $D = (d_0, \dots, d_n) \in \mathbb{Z}^{n+1}$, let us define $\deg^+(D) := \deg(D^+) = \sum_{i\,:\,d_i \geq 0} d_i$ and $\deg^-(D) := \deg(D^-) = \sum_{i\,:\,d_i \leq 0} d_i$. We have:

Lemma 1.2.9. *For every integral point $D \in \mathbb{Z}^{n+1}$,*

$$r(D) \;=\; \min\{\, \deg^+(\nu + D) \mid \nu \in \mathrm{Ext}(L)\,\} - 1\,.$$

Proof. First recall that

$$\begin{aligned} r(D) \;&=\; \min\{\, \deg(E) \mid |D - E| \;=\; \emptyset \quad \text{and} \quad E \geq 0\,\} - 1 \\ &=\; \min\{\, \deg(E) \mid E - D \in \Sigma(L) \text{ and } \quad E \geq 0\,\} - 1 \quad \text{(By Lemma 1.2.2).} \end{aligned}$$

Let $E \geq 0$ and $p = E - D$ be a point in $\Sigma(L)$. By Theorem 1.2.6, we know that p is a point in $\Sigma(L)$ if and only if $p = \nu + E'$ for some point ν in $\mathrm{Ext}(L)$ and $E' \geq 0$. So we can write $E = p + D = \nu + E' + D$ where $\nu \in \mathrm{Ext}(L)$ and $E' \geq 0$. Hence we have

$$r(D) = \min\{\, \deg(\nu + E' + D) \mid \nu \in \mathrm{Ext}(L),\ E' \geq 0 \text{ and } \nu + E' + D \geq 0\,\} - 1.$$

We now observe that for every $\nu \in \mathbb{Z}^{n+1}$, the integral point $E' \geq 0$ of minimum degree such that $E' + \nu + D \geq 0$ has degree exactly $\deg^+(-\nu - D)$. We infer that

$$\begin{aligned} \deg(\nu + E' + D) \;&=\; \deg(E') + \deg(\nu + D) \;=\; \deg^+(-\nu - D) + \deg(\nu + D) \\ &=\; \deg^-(\nu + D) + \deg(\nu + D) \;=\; \deg^+(\nu + D). \end{aligned}$$

We conclude that $r(D) = \min\{\deg^+(\nu + D) | \nu \in \mathrm{Ext}(L)\} - 1$, and the lemma follows. □

1.2.3 Min- and Max-Genus of SubLattices of A_n and Uniform Lattices

We define two notions of genus for full-rank sublattices of A_n, *min- and max-genus*, in terms of the extremal points of the Sigma-Region of L. (The same definition works for full-rank sublattices of H_0.)

Definition 1.2.10. *(Min- and Max-Genus) The min- and max-genus of a given sublattice L of A_n of rank n, denoted respectively by g_{min} and g_{max}, are defined as follows:*

$$\begin{aligned} g_{min}(L) &= \inf\{\, -\deg(\nu) \mid \nu \in \mathrm{Ext}(L)\,\} + 1\,. \\ g_{max}(L) &= \sup\{\, -\deg(\nu) \mid \nu \in \mathrm{Ext}(L)\,\} + 1\,. \end{aligned}$$

Remark 1.2.11. There are some other notions of genus associated to a given lattice, for example., the notion *spinor genus* for lattices developed by Eichler (see polyhedraEic52 and [31]) in the context of integral quadratic forms. Every sublattice of A_n provides a quadratic form in a natural way. But a priori there is no relation between these notions.

It is clear by definition that $g_{min} \leq g_{max}$. But generally these two numbers could be different.

Definition 1.2.12. *A sublattice $L \subseteq A_n$ of rank n is called uniform if $g_{min} = g_{max}$. The genus of a uniform sublattice is $g = g_{min} = g_{max}$.*

As we will show later in Chapter 2, sublattices generated by Laplacian of graphs are uniform.

1.3 Proofs of Theorem 1.2.6 and Theorem 1.2.7

In this section, we present the proofs of Theorem 1.2.6 and Theorem 1.2.7.

Recall that $\Sigma^{\mathbb{R}}(L)$ is the set of points in $\mathbb{R}^{n+1}$ that are not dominated by any point in L and $\Sigma^c(L)$ is the topological closure of $\Sigma^{\mathbb{R}}(L)$ in $\mathbb{R}^{n+1}$. Also, recall that $\mathrm{Ext}^c(L)$ denotes the set of extremal points of $\Sigma^c(L)$. These are the set of points which are local minimum of the degree function. As we said before, instead of working with the Sigma-Region directly, we initially work with $\Sigma^c(L)$. We first prove Theorem 1.2.7. Namely, we prove $\Sigma^c(L) = \cup_{\nu \in \mathrm{Ext}^c(L)} H^+_\nu$. To prepare for the proof of this theorem, we need a series of lemmas.

The following lemma provides a description of $\Sigma^c(L)$ in terms of the domination order in $\mathbb{R}^{n+1}$. Recall that for two points $x = (x_0, \ldots, x_n)$ and $y = (y_0, \ldots, y_n)$, $x \leq y$ (resp. $x < y$) if $x_i \leq y_i$ (resp. $x_i < y_i$) for all $0 \leq i \leq n$.

Lemma 1.3.1. $\Sigma^c(L) \;=\; \{\, p \mid p \in \mathbb{R}^{n+1}$ *and* $\;\forall\, q \in L \;:\; p \not< q\,\}$.

Proof. Let $\bar{\Sigma}^c(L) := \{p | p \in \mathbb{R}^n \text{and} \forall p' \in L : p \not< p'\}$. We write $\Sigma^c(L) = \Sigma^{\mathbb{R}}(L) \cup \partial\Sigma^{\mathbb{R}}(L)$. For a point p is in $\Sigma^{\mathbb{R}}(L)$, we have $p \not\leq p'$ for all p' in L and so $\Sigma^{\mathbb{R}} \subseteq \bar{\Sigma}^c(L)$. We now prove that $\partial\Sigma^{\mathbb{R}}(L) \subseteq \bar{\Sigma}^c(L)$. Let p be a point in $\partial\Sigma^{\mathbb{R}}(L)$ and suppose that p is not contained in $\Sigma^{\mathbb{R}}(L)$. We have $p \leq p'$ for some points p' in L, and that $p + \delta(1, \ldots, 1)$ is contained in $\Sigma^{\mathbb{R}}(L)$ for all $\delta > 0$. But this means that for any point p' in L such that $p \leq p'$, there exists some i such that $(p)_i = (p')_i$ and hence $p \not< p'$. This proves that $\partial\Sigma^c(L) \subseteq \bar{\Sigma}^c(L)$, and so $\Sigma^c(L) \subseteq \bar{\Sigma}^c(L)$. We now verify that $\bar{\Sigma}^c(L) \subseteq \Sigma^c(L)$. We conclude that $\Sigma^c(L) = \{ p \,|\, p \in \mathbb{R}^n \text{ and} \, \forall\, p' \in L \,:\, p \not< p' \}$. $\square$

Lemma 1.3.2. *Extremal points of $\Sigma^c(L)$ are contained in $\partial(\Sigma^c(L))$.*

Let p be a point in $\Sigma^c(L)$ and let d be a vector in $\mathbb{R}^{n+1}$. We say that d is *feasible* for p, if it satisfies the following properties:

1. $\deg(d) < 0$.
2. There exists a $\delta_0(p, d) > 0$ such that for every $0 \leq \delta \leq \delta_0(p, d)$, $p + \delta d \in \Sigma^c(L)$. By Lemma 1.3.1, this means that $p + \delta d \not< p'$ for all lattice points $p' \in L$.

Furthermore, we define the function $\epsilon_{p,d} \,:\, L \to \mathbb{R} \cup \{\infty\}$ as follows:

$$\epsilon_{p,d}(q) \;=\; \inf \{ \epsilon \,|\, \epsilon > 0 \text{ and } p + \epsilon d < q \}.$$

Let $I \;=\; \{ i \,|\, 0 \leq i \leq n \text{ and } p_i \geq q_i \}$. We have the following explicit description of $\epsilon_{p,d}$.

$$\epsilon_{p,d}(q) = \begin{cases} 0 & \text{if } I = \emptyset. \\ \max_{i \in I} \frac{(q_i - p_i)}{d_i} & \text{if } I \neq \emptyset,\ \forall\ i \in I,\ \ d_i < 0, \text{ and } \exists\ \epsilon > 0 \text{ such that } p + \epsilon d < q, \\ \infty & \text{otherwise.} \end{cases} \tag{1.2}$$

We verify that

Lemma 1.3.3. *For a point p in $\Sigma^c(L)$, $\epsilon_{d,p}(q) \;\geq\; \epsilon_{d^-,p}(q)$ for all $q \in L$. In the only cases when the inequality is strict, we must have $\epsilon_{d,p}(q) = \infty$ and $\epsilon_{d^-,p}(q) > 0$.*

We now prove the following lemma which links the function $\epsilon_{d,p}$ to the feasibility of d at p.

Lemma 1.3.4. *For a point p in $\Sigma^c(L)$ and d in $\mathbb{R}^{n+1}$ with $\deg(d) < 0$, d is not feasible for p if and only if $\epsilon_{p,d}(q) = 0$ for some $q \in L$.*

Proof. Let p be a point of $\Sigma^c(L)$.

($\Rightarrow$). Assume the contrary, then we should have the following properties:

1. $\deg(d) \;<\; 0$,
2. $\epsilon_{p,d}(q) > 0$ for all $q \in L$,

We claim that $\inf_{q \in L} \{ \epsilon_{p,d}(q) \} \;>\; \delta_0$, for some $\delta_0 \;>\; 0$. By the definition of $\epsilon_{p,d}$ and the fact that L is a sublattice of A_n we know that if $\epsilon_{p,d}(q) \neq 0$, then since L is a sublattice of A_n and hence every element of L has integer coordinates, we know that $\epsilon_{p,d}(q)$ is at least $\min_{\{i:\, d_i<0\}} \frac{\{p_i\}}{|d_i|}$, where $0 < \{p_i\} = p_i - \lceil p_i - 1 \rceil \leq 1$ is the rational part of p_i if p_i is not integral, and is 1 if p_i is integral. As the number of indices is finite, we conclude that $\delta_0 = \min_{\{i:\, d_i<0\}} |\frac{\{p_i\}}{d_i}|$ and the claim holds. It follows that $p + \epsilon d \not< q$ for all q in L and for all $0 \leq \epsilon \leq \delta_0$. This implies that d is feasible for p.

($\Leftarrow$). If $\epsilon_{p,d}(q) = 0$ for some $q \in L$, then there exists a $\delta_0 > 0$ such that $p + \delta d < p'$ for every $0 < \delta \leq \delta_0$. This shows that d is not feasible for p. □

Remark 5. Note that in the proof of Lemma 1.3.4 the argument that $\epsilon_{p,d}(q)$ is at least $\min_{\{i:\, d_i<0\}} \frac{\{p_i\}}{|d_i|}$ requires that L is a sublattice of A_n. The argument can be modified to handle arbitrary sublattices of H_0 by using the fact that a lattice is a discrete subset of H_0. The argument is as follows: fix an $\epsilon_0 > 0$ and observe that there can only a finite number of lattice points S in L that satisfy $\epsilon(p,d)(q) \leq \epsilon_0$. We can now pick the minimum $\epsilon(p,d)(q)$ among points in S to conclude that $\inf_{q \in L} \{ \epsilon_{p,d}(q) \} > 0$.

Corollary 1.3.5. *For a point p in $\Sigma^c(L)$, p is an extremal point if and only if for every vector $d \in \mathbb{R}^{n+1}$ with $\deg(d) < 0$, we have $\epsilon_{p,d}(q) = 0$ for some q in L.*

Combining Lemma 1.3.3 and Corollary 1.3.5, we obtain the following result:

Lemma 1.3.6. *If p is not an extremal point of $\Sigma^c(L)$, then there exists a vector d in H_O^- which is feasible for p.*

Proof. If p is not an extremal point of $\Sigma^c(L)$, then there exists a vector d_0 in $\mathbb{R}^{n+1}$ that is feasible for p. By Corollary 1.3.5, d_0 has the following properties:

1. $\deg(d_0) < 0$,

2. $\epsilon_{d_0,p}(q) > 0$ for all $q \in L$,

Let $d := d_0^-$. We have $\deg(d) < 0$, since $\deg(d_0) < 0$ and $d = d_0^-$. By Lemma 1.3.3, we have $\epsilon_{d_0,p}(q) \geq \epsilon_{d,p}(q)$ for all $q \in L$, and in the only cases for q when the inequality is strict we have $\epsilon_{d,p}(q) > 0$. We infer that d also satisfies Properties 1 and 2. By Corollary 1.3.5, d is also feasible for p and by construction, d belongs to H_O^-; the lemma follows. □

Consider the set $\deg(\Sigma^c(L)) = \{ \deg(p) \mid p \in \Sigma^c(L) \}$. The next lemma shows that the degree function is bounded below on the elements of $\Sigma^c(L)$ (by some negative real number).

Lemma 1.3.7. *For a rank n sublattice L of A_n, $\inf(\deg(\Sigma^c(L))$ is finite.*

Proof. It is possible to give a direct proof of this lemma. But using our results in Section 1.4 allows us to shorten the proof. So we postpone the proof to Section 1.4. □

We are now in a position to present the proofs of Theorem 1.2.7 and Theorem 1.2.6.

Proof of Theorem 1.2.7. Consider a point p in $\Sigma^c(L)$. We should prove the existence of an extremal point $\nu \in \mathrm{Ext}^c(L)$ such that $\nu \leq p$.

Consider the cone H_p^-. As a consequence of Lemma 1.3.7, we infer that the region $\Sigma^c(L) \cap H_p^-$ is a bounded closed subspace of $\mathbb{R}^{n+1}$, and so it is compact. The degree function deg restricted to this compact set, achieves its minimum on some point $\nu \in \Sigma^c(L) \cap H_p^-$. We claim that $\nu \in \mathrm{Ext}^c(L)$. Suppose that this is not the case. By Lemma 1.3.6, there exists a feasible vector $d \in H_O^-$ for ν, i.e., such that $\nu + \delta d \in \Sigma^c(L)$ for all sufficiently small $\delta > 0$. Now we verify that

- $\nu + \delta\, d \in H_p^-$ and hence $\nu + \delta d \leq p$,
- $\deg(\nu + \delta\, d) < \deg(\nu)$.

This contradicts the choice of ν. □

Proof of Theorem 1.2.6. In order to establish Theorem 1.2.6, we first prove that every point in $\mathrm{Ext}^c(L)$ is an integral point. For the sake of a contradiction, suppose that there exists a non integral point in $\mathrm{Ext}^c(L)$. Let $p = (p_0, \ldots, p_n)$ be such a point and suppose without loss of generality that p_0 is not integer. We claim that the vector $d = -e_0 = (-1, 0, 0, \ldots, 0)$ is feasible. We have $\epsilon_{p,d}(q) > 0$ for all $q \in L$, and so by Corollary 1.3.5 we conclude that p could not be an extremal point of $\Sigma^c(L)$.

Let $\Sigma^c_{\mathbb{Z}}(L)$ be the set of integral points of $\Sigma^c(L)$. We show that $\Sigma^c_{\mathbb{Z}}(L)+(1,\ldots,1) = \Sigma(L)$. Note that as soon as this is proved, Theorem 1.2.7 and the fact that extremal points of $\Sigma^c(L)$ are all integral points implies Theorem 1.2.6.

We prove $\Sigma^c_{\mathbb{Z}}(L) + (1,\ldots,1) \subseteq \Sigma(L)$.— Let $u = v + (1,\ldots,1) \in \Sigma^c_{\mathbb{Z}}(L) + (1,\ldots,1)$, for a point $v \in \Sigma^c_{\mathbb{Z}}(L)$. To show $u \in \Sigma(L)$ we should prove that $\forall q \in L : u \nleq q$. Suppose that this is not the case and let $q \in L$ be such that $u \leq q$. It follows that $u - (1,\ldots,1) < q$ and hence, $v \notin \Sigma^c(L)$, which is a contradiction.

We prove $\Sigma(L) \subseteq \Sigma^c_{\mathbb{Z}}(L) + (1,\ldots,1)$.— A point u in $\partial\Sigma^c(L)$ is contained in $H^-(q)$ for some q in L and hence $u \leq q$. We infer that $\Sigma(L)$ is contained in the interior of $\Sigma^c(L)$, and so for each point p of $\Sigma(L)$, every vector in $\mathbb{R}^{n+1}$ of negative degree will be feasible. By Lemma 1.3.7, there exists a point $p_c \in \partial\Sigma^c(L)$ such that $p = p_c + t(1,\ldots,1)$ for some $t > 0$. It follows that $p > p_c$. By Theorem 1.2.7, $p_c \in H_\nu^+$ for some ν in $\mathrm{Ext}^c(L)$. This implies that $p > \nu$ for some $\nu \in \mathrm{Ext}^c(L)$. By definition, p is an integral point and we just showed that ν is also an integral point. Hence we can further deduce that $p \geq \nu + (1,\ldots,1)$. We infer that $p - (1,\ldots,1) \geq \nu$ and therefore, $p - (1,\ldots,1) \in \Sigma^c(L)$ (because $H_\nu^+ \subset \Sigma^c(L)$). It follows that $p \in \Sigma^c_{\mathbb{Z}}(L) + (1,\ldots,1)$.

We finally show that every $\nu \in \mathrm{Ext}^c(L) + (1,\ldots,1)$ must be have minimum degree among all points in $N(\nu) \cap \Sigma(L)$, the set of integer points that are at distance at most one in the ℓ_1-norm since, otherwise suppose that there a point in q in $N(\nu) \cap \Sigma(L)$ with $\deg(q) < \deg(\nu)$. Hence, we know that $q \leq \nu$. Now, consider a point $\nu' \in \mathrm{Ext}(L)$ such that $q \geq \nu'$, this means that $\nu \geq \nu'$. But, this contradicts our assumption that $\nu \in \mathrm{Ext}(L)$. □

1.4 Lattices under a Simplicial Distance Function

In this section, we provide some basic properties of the Voronoi diagram of a sublattice L of A_n under a simplicial distance function $d_\triangle(.,.)$ which we define below. The distance function $d_\triangle(.,.)$ has the following explicit form, and as we will see in this section, is the distance function having the homotheties of the standard simplex in H_0 as its balls (which explains the name simplicial distance function). For two points p and q in H_0, the simplicial distance between p and q is defined as follows

$$d_\triangle(p,q) := \inf\Big\{\lambda \,|\, q - p + \lambda(1,\dots,1) \geq 0\Big\}.$$

The basic properties of $d_\triangle$ are better explained in the more general context of polyhedral distance functions that we now explain.

1.4.1 Polyhedral Distance Functions and their Voronoi Diagrams

Let Q be a convex polytope in $\mathbb{R}^n$ with the reference point $O = (0,\dots,0)$ in its interior. The *polyhedral distance function* $d_Q(.,.)$ between the points of $\mathbb{R}^n$ is defined as follows:

$$\forall\, p,q \in \mathbb{R}^n,\; d_Q(p,q) \;:=\; \inf\{\lambda \geq 0 \,|\, q \in p + \lambda.Q\}, \text{ where } \quad \lambda.Q \;=\; \{\, \lambda.x \,|\, x \in Q \,\}.$$

d_Q is not generally symmetric, indeed it is easy to check that $d_Q(.,.)$ is symmetric if and only if the polyhedron Q is centrally symmetric i.e., $Q = -Q$. Nevertheless $d_Q(.,.)$ satisfies the triangle inequality.

Lemma 1.4.1. *For every three points $p,q,r \in \mathbb{R}^n$, we have $d_Q(p,q)+d_Q(q,r) \geq d_Q(p,r)$. In addition, if q is a convex combination of p and r, then $d_Q(p,q) + d_Q(q,r) = d_Q(p,r)$.*

Proof. To prove the triangle inequality, it will be sufficient to show that if $q \in p + \lambda.Q$ and $r \in q + \mu.Q$, then $r \in p + (\lambda+\mu).Q$. We write $q = p + \lambda.q'$ and $r = q + \mu.r'$ for two points q' and r' in Q. We can then write $r = p+\lambda.q'+\mu.r' = p+(\lambda+\mu)(\frac{\lambda}{\lambda+\mu}.q' + \frac{\mu}{\lambda+\mu}.r')$. Q being convex and $\lambda,\mu \geq 0$, we infer that $\frac{\lambda}{\lambda+\mu}.q' + \frac{\mu}{\lambda+\mu}.r' \in Q$, and so $r \in p+(\lambda+\mu).Q$. The triangle inequality follows.

To prove the second part of the lemma, let $t \in [0,1]$ be such that $q = t.p + (1-t).r$. By the triangle inequality, it will be enough to prove that $d_Q(p,q) + d_Q(q,r) \leq d_Q(p,r)$. Let $d_Q(p,r) = \lambda$ so that $r = p + \lambda.r'$ for some point r' in Q. We infer first that $q = t.p + (1-t).r = t.p + (1-t)(p + \lambda.r') = p + (1-t)\lambda.r'$, which implies that $d_Q(p,q) \leq (1-t)\lambda$. Similarly we have $t.r = t.p + t\lambda.r' = q - (1-t)r + t\lambda.r'$. It follows that $r = q + t\lambda r'$ and so $d_Q(q,r) \leq t\lambda$. We conclude that $d_Q(p,q) + d_Q(q,r) \leq d_Q(p,r)$, and the lemma follows. $\square$

We also observe that the polyhedral metric $d_Q(.,.)$ is translation invariant, i.e.,

Lemma 1.4.2. *For any two points p,q in $\mathbb{R}^n$, and for any vector $v \in \mathbb{R}^n$, we have $d_Q(p,q) = d_Q(p-v,q-v)$. In particular, $d_Q(p,q) = d_Q(p-q,O) = d_Q(O,q-p)$.*

Proof. If $q \in p + \lambda.Q$, then $q - v \in p - v + \lambda.Q$, and vice versa. $\square$

Remark 1.4.3. The notion of a polyhedral distance function is essentially the concept of a gauge function of a convex body that has been studied in [77]. Lemmas 1.4.1 and 1.4.2 can be derived in a straight forward way from the results in [77]. We include them here for the sake of reference.

Consider a discrete subset $\mathcal{S}$ in $\mathbb{R}^n$. For a point s in $\mathcal{S}$, we define the *Voronoi cell* of s with respect to d_Q as $V_Q(s) = \{p \in \mathbb{R}^n \,|\, d_Q(p, s) \leq d_Q(p, s') \text{ for any other point } s' \in \mathcal{S}\}$. The *Voronoi diagram* $\mathrm{Vor}_Q(\mathcal{S})$ is the decomposition of $\mathbb{R}^n$ induced by the cells $V_Q(s)$, for $s \in \mathcal{S}$. We note however that this need not be a cell decomposition in the usual sense.

We state the following lemma on the shape of cells $V_Q(s)$.

Lemma 1.4.4. *[25] Let $\mathcal{S}$ be a discrete subset of $\mathbb{R}^n$ and $\mathrm{Vor}_Q(\mathcal{S})$ be the Voronoi cell decomposition of $\mathbb{R}^n$. For any point s in $\mathcal{S}$, the Voronoi cell $V_Q(s)$ is a star-shaped set with s as a kernel.*

Proof. Assume the contrary. Then there is a line segment $[s, r]$ and a point q between s and r such that $r \in V_Q(s)$ and $q \notin V_Q(s)$. Suppose that q is contained in $V(s')$ for some $s' \neq s$. We should then have $d_Q(q, s) > d_Q(q, s')$. By Lemma 1.4.1, $d_Q(r, s) = d_Q(r, q) + d_Q(q, s)$. We infer that

$$d_Q(r, s) \;=\; d_Q(r, q) + d_Q(q, s) > d_Q(r, q) + d_Q(q, s') \geq \; d_Q(r, s'), \text{ contradicting } r \in V_Q(s).$$

$\square$

Definition 1.4.5. (Voronoi Neighbours) *We say that distinct points p, $q \in S$ are Voronoi neighbours if the intersection of their Voronoi cells is non-empty.*

1.4.2 Voronoi Diagram of SubLattices of A_n

Voronoi diagrams of root lattices under the Euclidean metric have been studied previously in literature. Conway and Sloane [30, 31], describe the Voronoi cell structure of root lattices and their duals under the Euclidean metric.

Here we study Voronoi diagrams of sublattices of A_n under polyhedral distance functions (and later under the simplicial distance functions $d_{\triangle}(. \, , .)$). We will see the importance of this study in the proof of Riemann-Roch Theorem in Section 1.5, and in the geometric study of the Laplacian of graphs in Chapter 2.

Let L be a sublattice of A_n of full rank. Note that L is a discrete subset of the hyperplane H_0 and $H_0 \simeq \mathbb{R}^n$. Let $Q \subset H_0$ be a convex polytope of dimension n in H_0. We will be interested in the Voronoi cell decomposition of the hyperplane H_0 under the distance function $d_Q(., .)$ induced by the points of L. The following lemma, which essentially uses the translation-invariance of $d_Q(. \, , .)$, shows that these cells are all simply translations of each other.

Lemma 1.4.6. *For a point p in $L, V_Q(p) = V_Q(O) + p$. As a consequence,* $\mathrm{Vor}_Q(L) = V_Q(O) + L$.

Proof. Let p_1 and p_2 be two points of L, and $q_1 \in V_Q(p_1)$ be a point of the Voronoi cell of p_1. By definition, we should have $d_Q(q_1, p_1) \leq d(q_1, p)$ for every point p in L. Let $q_2 := q_1 - p_1 + p_2$. By the translation invariance of d_Q, we have $d_Q(q_2, p_2) = d_Q(q_1, p_1)$. We claim that $q_2 \in V_Q(p_2)$. We should prove that $d_Q(q_2, p_2) \leq d_Q(q_2, p)$ for every point $p \in L$. Let p be in L. As L is a lattice, and $p_1, p_2 \in L$, the point $p - p_2 + p_1$ is also in L. We infer that

$$d_Q(q_2, p_2) \ = \ d_Q(q_1, p_1) \leq d_Q(q_1, p - p_2 + p_1) = d_Q(q_2, p),$$

and the claim follows. Note that the last equality above comes again from the translation-invariance of d_Q. The above arguments show that $V_Q(p_1) + p_2 - p_1 \subseteq V_Q(p_2)$. Replacing the role of p_2 by p_1, we finally infer that $V_Q(p_2) = V_Q(p_1) + p_2 - p_1$. For $p_1 = O$ and $p_2 = p \in L$, we obtain $V(p) = V(O) + p$. □

By Lemma 5.1.7, to understand the Voronoi cell decomposition of H_0, it will be enough to understand the cell $V_Q(O)$. We already know that $V_Q(O)$ is a star-shaped set. The following lemma shows that $V_Q(O)$ is compact, and so it is a (not necessarily convex) star-shaped polytope.

Lemma 1.4.7. *The Voronoi cell $V_Q(O)$ is compact.*

Proof. It is sufficient to prove that $V_Q(O)$ does not contain any infinite ray. Indeed, $V_Q(O)$ being star-shaped and closed, this will imply that $V_Q(O)$ is bounded and so we have the compactness.

Assume, for the sake of a contradiction, that there exists a vector $v \neq O$ in H_0 such that the ray $t.v$ for $t \geq 0$ is contained in $V_Q(O)$. This means that

$$\text{For every } t \geq 0 \text{ and for every } p \in L, \text{ we have } d_Q(t.v, O) \leq d_Q(t.v, p). \qquad (1.3)$$

Choose a real number λ such that $0 < \lambda < d_Q(v, O)$. By Lemma 1.4.1, $d_Q(t.v, O) = td_Q(v, O) > \lambda t$ for $t > 0$. By the definition of d_Q, the choice of λ and Property (1.3), the polytope $t.v + t\lambda.Q = t.(v + \lambda Q)$ does not contain any point $p \in L$ for $t > 0$. Let $\mathcal{C} = \bigcup_{t \geq 0} t.(v + \lambda.Q)$. We verify that $\mathcal{C}$ is the cone generated by $v + \lambda.Q$. It follows that $\mathcal{C}$ does not contain any lattice point apart from O (for $t = 0$). In addition, Q being a polytope of dimension n, $\mathcal{C}$ should be a cone of full dimension in H_0. But this will provide a contradiction, because as we will show below for any vector $\bar{v}$ with rational coordinates in H_0, the open ray $t.\bar{v}$ for $t > 0$ contains a lattice point in L. (And since, rational numbers are dense in the space of real numbers, we know that any cone $\mathcal{C}$ of full dimension in H_0 contains a rational vector.) To see this, observe that a basis for L is also a basis for the n-dimensional $\mathbb{Q}$-vector space $H_0(\mathbb{Q})$. Here $H_0(\mathbb{Q})$ denotes the rational points of the hyperplane H_0. This means that $\bar{v}$ can be written as a rational combination of some points in L. Multiplying by a sufficiently large integer number N, $N.\bar{v}$ can be written as an integral combination of the same points in L, i.e., $N.\bar{v} \in L$, and this finishes the proof of the lemma. □

Lemma 1.4.8. *Any shortest vector of a lattice L under the polyhedral distance function d_Q is a Voronoi neighbour of the origin under the same distance function.*

Proof. Let p be a shortest vector of L in the distance function d_Q. Consider the intersection p_I of the ray $\overrightarrow{Op}$ with the Voronoi cell $V_Q(O)$ of the origin under the distance function d_Q and by Lemma 1.4.7, we know that the ray $\overrightarrow{Op}$ intersects $V_Q(O)$ at some point $p_I \neq p$. By Lemma 1.4.1, we have: $d_Q(O,p) = d_Q(O,p_I) + d_Q(p_I,p)$. Assume for contradiction that p is not a Voronoi neighbour of the origin under the distance function d_Q. This means that there is a lattice point p' such that $d_Q(p_I,p') < d_Q(p_I,p)$. Hence we have:

$$d_Q(O,p') \leq d_Q(O,p_I) + d_Q(p_I,p') < d_Q(O,p_I) + d_Q(p_I,p) = d_Q(O,p). \tag{1.4}$$

This contradicts our assumption that p is a shortest vector of L. $\square$

We will mainly be interested in two special polytopes $\triangle$ and $\bar{\triangle}$ in H_0. They are both standard simplices of H_0 under an appropriate isometry $H_0 \simeq \mathbb{R}^n$. The n-dimensional regular simplex $\triangle(O)$ centered at the origin O has vertices at the points $b_0, b_1, \ldots, b_n$. For all $0 \leq i,j \leq n$, the coordinates of b_i are given by:

$$(b_i)_j = \begin{cases} n & \text{if i=j,} \\ -1 & \text{otherwise.} \end{cases}$$

The simplex $\bar{\triangle}(O)$ is the *opposite* simplex to $\triangle(O)$, i.e., $\bar{\triangle}(O) := -\triangle(O)$. The simplicial distance functions $d_\triangle(.\,,.)$ and $d_{\bar{\triangle}}(.\,,.)$ are the distance functions in $\mathbb{R}^{n+1}$ defined by $\triangle$ and $\bar{\triangle}$ respectively. It is easy to check the following anti-symmetric property for the above distance functions: *For any pair of points $p,q \in \mathbb{R}^{n+1}$, we have $d_\triangle(p,q) = d_{\bar{\triangle}}(q,p)$.* (This is indeed true for any convex polytope Q: $d_Q(p,q) = d_{\bar{Q}}(q,p)$, where $\bar{Q} = -Q$.)

1.4.3 Geometric Invariants of a Lattice with respect to Polyhedral Distance Functions

We will define some important invariants of a lattice with respect to polyhedral distance functions.

Definition 1.4.9. *An element q of L is called a shortest vector with respect to the polyhedral distance function $d_\mathcal{P}$ if $d_\mathcal{P}(O,q) \leq d_\mathcal{P}(O,q')$ for all $q' \in L/\{O\}$, where O is the origin. We denote $d_\mathcal{P}(O,q)$ by $\nu_\mathcal{P}(L)$.*

Note that the shortest vector with respect the distance functions $d_\mathcal{P}$ and $d_{\bar{\mathcal{P}}}$ can be potentially different.

Definition 1.4.10. *For a lattice L, we define packing and covering radius of L with respect $\mathcal{P}$ as:*

$Pac_\mathcal{P}(L) = \sup\{R|\ \mathcal{P}(q_1,R) \cap \mathcal{P}(q_2,R) = \emptyset,\ \forall q_1,\ q_2 \in L\ q_1 \neq q_2\}$

$Cov_\mathcal{P}(L) = \inf\{R|\ \ \textit{every } p \in Span(L) \textit{ is contained in } \mathcal{P}(q,R) \textit{ for some } q \in L\}$

Note that for polyhedral distance functions it is no longer true that $\mathrm{Pac}_{\mathcal{P}}(L) = \nu_{\mathcal{P}}(L)/2$. We will see examples in Chapter 3

Notation. In the following we will use the following terminology: For a point $v \in H_0$, we let $\triangle(v) = v + \triangle(O)$ and $\bar{\triangle}(v) = v + \bar{\triangle}(O)$. More generally given a real $\lambda \geq 0$ and $v \in H_0$, we define $\triangle_\lambda(v) = v + \lambda \cdot \triangle(O)$, and similarly, $\bar{\triangle}_\lambda(v) = v + \lambda \cdot \bar{\triangle}(O)$. We can think of these as *balls of radius* λ around v for $d_\triangle$ and $d_{\bar{\triangle}}$ respectively.

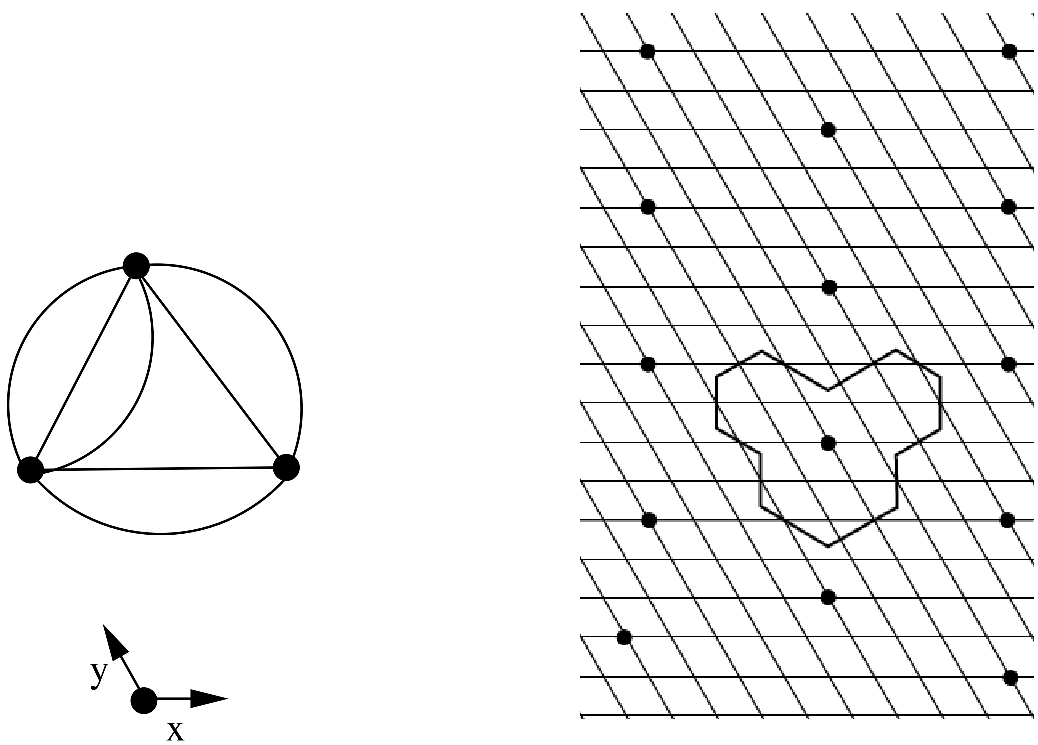

Figure 1.2: The shape of a Voronoi-cell in the Laplacian lattice of a graph with three vertices. The multi-graph G has three vertices and 7 edges. The lattice A_2 is generated by the two vectors $x = (1, -1, 0)$ and $y = (-1, 0, 1)$. The corresponding Laplacian sublattice of A_2, whose elements are denoted by $\bullet$, is generated by the vectors $(-5, 3, 2) = -3x + 2y$ and $(3, -5, 2) = 5x + 2y$ (and $(2, 2, -4) = -2x - 4y$), which correspond to the vertices of G.

The following lemma shows that the definition given in the beginning of this section coincides with the definition of $d_\triangle$ given above. We can explicitly write a formula for $d_\triangle(.,.)$ and $d_{\bar{\triangle}}(.,.)$ in the hyperplane H_0:

Lemma 1.4.11. *For two points $p = (p_0, p_1, \ldots, p_n)$ and $q = (q_0, q_1, \ldots, q_n)$ in H_0, the $\triangle$-simplicial distance from p to q is given by $d_\triangle(p, q) = |\bigoplus_{i=0}^n (q_i - p_i)|$. And the $\bar{\triangle}$-simplicial distance from p to q is given by $d_{\bar{\triangle}}(p, q) = |\bigoplus_{i=0}^n (p_i - q_i)|$. Here the sum $\bigoplus_i (x_i - y_i)$ denotes the tropical sum of the numbers $x_i - y_i$.*

Proof. By the anti-symmetry property of the distance function $d_\triangle(.,.)$ (namely $d_\triangle(p,q) = d_{\bar{\triangle}}(q,p),\ \forall p,q$), we only need to prove the lemma for $d_\triangle(.\,,.)$. By definition, $d_\triangle(p,q)$ is the smallest positive real λ such that $q \in p + \lambda.\triangle$. The simplex $\triangle$ being the convex hull of the vectors b_i defined above, it follows that for an element $x \in \lambda.\triangle$, there should exist non-negative reals $\mu_i \geq 0$ such that $\sum_{i=0}^n \mu_i = \lambda$ and $x = \mu_0 b_0 + \mu_1 b_1 + \cdots + \mu_n b_n$. From the definition of the vector b_i's, we obtain $x = (n+1)(\mu_0, \mu_1, \ldots, \mu_n) - \lambda(1,\ldots,1)$. It follows that $d_\triangle(p,q)$ is the smallest λ such that $q - p + \lambda.(1,\ldots,1)$ becomes equal to $(n+1)(\mu_0,\mu_1,\ldots,\mu_n)$ for some $\mu_i \geq 0$ such that $\sum_i \mu_i = \lambda$. Let λ_0 be the smallest positive real number such that the vector $\mu := \frac{1}{n+1}(q - p + \lambda_0.(1,\ldots,1))$ has non-negative coordinates. As $p,q \in H_0$, a simple calculation shows that the other condition $\sum_i \mu_i = \lambda_0$ holds automatically, and hence such λ_0 is equal to $d_\triangle(p,q)$. By construction $\lambda_0 = \max_i (p_i - q_i) = -\min_i (q_i - p_i)$. It follows that $d_\triangle(p,q) = |\bigoplus_{i=0}^n (q_i - p_i)|$. $\square$

1.4.4 Vertices of $\mathrm{Vor}_\triangle(L)$ that are Critical Points of a Distance Function.

For a discrete subset $\mathcal{S}$ of H_0 (for example., $\mathcal{S} = L$), the simplicial distance function $h_{\triangle,\mathcal{S}} : H_0 \to \mathbb{R}$ is defined as follows:

$$h_{\triangle,\mathcal{S}}(x) \ = \ \bigoplus_{p\in\mathcal{S}} d_\triangle(x,p) \ = \ \min_{p\in\mathcal{S}} d_\triangle(x,p).$$

By definition, it is straightforward to verify that $h_{\triangle,\mathcal{S}}(x) = \sup\{\lambda \mid (x+\lambda\cdot\triangle)\cap\mathcal{S} = \emptyset\}$. Note that our definition above exactly imitates the classical definition of a distance function [40]. In what follows, we restrict ourselves to $\mathcal{S} = L$.

Remark 6. The notion of a simplicial distance function with respect to a lattice is sometimes captured in the language of "lattice point free simplices", see [74] for more details on this viewpoint. The author is indebted to Bernd Sturmfels for pointing out this connection. In joint work with Bernd Sturmfels is currently investigating this topic.

Let L be a full-rank sublattice of A_n and $h_{\triangle,L}$ be the distance function defined by L. We first give a description of $\partial\Sigma^c(L)$ (see Section 1.2.2) in terms of $h_{\triangle,L}$. The *lower-graph* of $h_{\triangle,L}$ is the graph of the function $h_{\triangle,L}$ in the negative half-space of $\mathbb{R}^{n+1}$, i.e., in the half-space of $\mathbb{R}^{n+1}$ consisting of points of negative degree. More precisely, the *lower-graph* of $h_{\triangle,L}$, denoted by $\mathrm{Gr}(h_{\triangle,L})$, consists of all the points $y - h_{\triangle,L}(y)(1,\ldots,1)$ for $y \in H_0$.

Lemma 1.4.12. *The lower-graph of $h_{\triangle,L}$ and $\partial\Sigma^c(L)$ coincide, i.e.,* $\mathrm{Gr}(h_{\triangle,L}) = \partial\Sigma^c(L)$.

In order to present the proof of Lemma 1.4.12, we need to make some remarks. Let p be a point of L. The function $f_p : H_0 \to \mathbb{R}^{n+1}$ is defined as follows:

$$\forall\, y \in H_0,\ f_p(y) := \sup\,\{y_t \mid y_t = y - t.(1,\ldots,1), t \geq 0, \text{ and } y_t \leq p\,\}.$$

Note that sup is defined with respect to the ordering of $\mathbb{R}^{n+1}$, and is well-defined because $y_t \geq y_{t'}$ if and only if $t \leq t'$. Remark also that $f_p(y)$ is finite.

Remark 1.4.13. The above notion has the following tropical meaning: Let $\lambda_p = \min\{t \in \mathbb{R} \mid t \odot p \oplus y = y\}$. Then $y_p = (-\lambda_p) \odot y$. The numbers λ_p are used in [33] to define the tropical closest point projection into some tropical polytopes. For a finite set of points $p_1, \ldots, p_l$ with the tropical convex-hull polytope Q, the tropical projection map π_Q at the point y is defined as $\pi_Q(y) = \lambda_{p_1} \odot p_1 \oplus \cdots \oplus \lambda_{p_l} \odot p_l$. It would be interesting to explore the connection between the work presented here and the theory of tropical polytopes.

A simple calculation shows that $f_p(y) = y - |\bigoplus_i (p_i - y_i)|.(1, \ldots, 1)$, and hence by Lemma 1.4.11, we obtain $f_p(y) = y - d_\triangle(y, p).(1, \ldots, 1)$. In other words, $f_p(y)$ is the lower-graph of the function $d_\triangle(., p)$. We claim that for all $y \in H_0$, $y - h_{\triangle,L}(y)(1, \ldots, 1) = \sup_{p\in L} f_p(y)$. Here, sup is understood as before with respect to the ordering of $\mathbb{R}^{n+1}$. In other words, the lower-graph $\mathrm{Gr}(h_{\triangle,L})$ is the lower envelope of the graphs $\mathrm{Gr}(f_p)$ for $p \in L$. To see this, remark that $\sup_{p\in L} f_p(y) = \sup_{p\in L}(y - d_\triangle(y, p).(1, \ldots, 1)) = y - (\min_{p\in L} d_\triangle(y, p)).(1, \ldots, 1) = y - h_{\triangle,L}(y).(1, \ldots, 1)$.

Proof of Lemma 1.4.12. By construction, for every point $y \in H_0$, the intersection of the half-ray $\{y - t(1, \ldots, 1) | t \geq 0\}$ with $\partial\Sigma^c(L)$ is the point $y - h_{\triangle,L}(L).(1, \ldots, 1) \in \mathrm{Gr}(h_{\triangle,L})$. More precisely, by the definition of $\Sigma^c(L)$ (see Section 1.2.2), we have

$$\begin{aligned} \partial\Sigma^c(L) &= \{\, z \mid z \leq p \text{ for some } p \in L \text{ and } z \not< p,\ \forall p \in L\} \\ &= \{\, \sup_{p\in L} f_p(y) \mid y \in H_0 \,\} \;=\; \mathrm{Gr}(h_{\triangle,L}) \quad \text{(By the discussion above).} \end{aligned}$$

□

It is possible to strengthen Lemma 1.4.12 and to obtain a description of the Voronoi diagram $\mathrm{Vor}_\triangle(L)$ in terms of the boundary of the Sigma-Region. The following lemma can be seen as the simplicial Voronoi diagram analogue of the classical result that *the Voronoi diagram under the Euclidean metric is the orthogonal projection of a lower envelope of paraboloids [35].*

Lemma 1.4.14. *The Voronoi diagram of L under the simplicial distance function $d_\triangle(.,.)$ is the projection of $\partial\Sigma^c(L)$ along $(1, \ldots, 1)$ onto the hyperplane H_0. More precisely, for any $p \in L$, the Voronoi cell $V_\triangle(p)$ is obtained as the image of $H_p^- \cap \partial\Sigma^c(L)$ under the projection map π_0.*

Proof. By definition, H_p^- consists of the points which are dominated by p. It follows that the intersection $H_p^- \cap \partial\Sigma^c(L)$ consists of all the points of $\partial\Sigma^c(L)$ which are dominated by p. By Lemma 1.4.12, the boundary of $\Sigma^c(L)$, $\partial\Sigma^c(L)$ coincides with the graph of the simplicial distance function $h_{\triangle,L}$. It follows that the intersection $H_p^- \cap \partial\Sigma^c(L)$ consists of all the points of the lower-graph of $h_{\triangle,L}$ that are dominated by p. By definition, any point of the lower-graph of $h_{\triangle,L}$ is of the form $y - h_{\triangle,L}(y).(1, \ldots, 1)$ for some $y \in H_0$. By definition of the function f_p, such a point is dominated by p if and only if $h_{\triangle,L}(y) \geq f_p(y)$. By definition, we know that $h_{\triangle,L}(y) \leq f_p(y)$ for all $y \in H_0$. We infer that for $y \in H_0$, $y - h_{\triangle,L}(y).(1, \ldots, 1) \in H_p^- \cap \partial\Sigma^c(L)$ if and only if $h_{\triangle,L}(y) = f_p(y)$, or equivalently, if and only if $y \in V_\triangle(p)$. We conclude that $V_\triangle(p) = \pi_0(H_p^- \cap \partial\Sigma^c(L))$ and this completes the proof of the lemma. □

As we show in the next two lemmas, it is possible to describe Voronoi *vertices* that are local maxima of $h_{\triangle,L}$ as the projection of the extremal points of the Sigma-Region onto the hyperplane H_0 (see below, Lemma 1.4.17, for a precise statement).

Let us denote by $\mathrm{Crit}(L)$ the set of all local maxima of $h_{\triangle,P}$. (In the example given in Figure 1.2, these are all the vertices of the polygon drawn in the plane H_2 (the right figure) having one concave and one convex neighbours on the polygon. There are six of them.)

Lemma 1.4.15. *The critical points of L are the projection of the extremal points of $\Sigma^c(L)$ along the vector* $(1,\ldots,1)$. *In other words,* $\mathrm{Crit}(L) = \pi_0(\mathrm{Ext}^c(L))$.

Proof. Let c be a point in $\mathrm{Crit}(L)$, and let $x = c - h_{\triangle,L}(c).(1,\ldots,1)$, be the corresponding point of the lower-graph of $h_{\triangle,L}$, $\mathrm{Gr}(h_{\triangle,L})$ ($= \partial\Sigma^c(L)$ by Lemma 1.4.12). Note that $\pi_0(x) = c$. We claim that $x \in \mathrm{Ext}^c(L)$. Assume the contrary. Then there should exist an infinite sequence $\{\, x_i \,\}_{i=1}^{\infty}$ such that (i) $x_i \in \partial\Sigma^c(L)$, (ii) $\deg(x_i) < \deg(x)$, and (iii) $\lim_{i\to\infty} x_i = x$. By (i) and Lemma 1.4.12, we can write $x_i = p_i - h_{\triangle,L}(p_i).(1,\ldots,1)$ for some $p_i \in H_0$. By (ii), we should have $-(n+1)h_{\triangle,L}(p_i) = \deg(x_i) < \deg(x) = -(n+1)h_{\triangle,L}(c)$ for every i, and so $h_{\triangle,L}(p_i) > h_{\triangle,L}(p_i)$. By (iii) and by the continuity of the map π_0, we have $\lim_{i\to\infty} p_i = c$. All together, we have obtained an infinite sequence of points $\{p_i\}$ in H_0 such that $h_{\triangle,L}(p_i) > h_{\triangle,L}(c)$ and $\lim_{i\to\infty} p_i = c$. This is a contradiction to our assumption that $c \in \mathrm{Crit}(L)$ is a local maximum of $h_{\triangle,L}$. A similar argument shows that for every point $x \in \mathrm{Ext}^c(L)$, $\pi_0(x)$ is in $\mathrm{Crit}(L)$, and the lemma follows. □

By Proposition 1.2.8, we have $\pi_0(\mathrm{Ext}^c(L)) = \pi_0(\mathrm{Ext}(L))$, and so

Corollary 1.4.16. *We have* $\mathrm{Crit}(L) = \pi_0(\mathrm{Ext}(L))$.

The following lemma gives a precise meaning to our claim that the critical points are the Voronoi vertices of the Voronoi diagram, and will be used in Chapter 2 in the proof of Theorem 2.1.9.

Lemma 1.4.17. *Each $v \in \mathrm{Crit}(L)$ is a vertex of the Voronoi diagram* $\mathrm{Vor}_{\triangle}(L)$*: there exist $n+1$ different points $p_0,\ldots,p_n$ in L such that $v \in \bigcap_i V(p_i)$. More precisely, a point $v \in H_0$ is critical, i.e., $v \in \mathrm{Crit}(L)$, if and only if it satisfies the following property: for each of the $n+1$ facets F_i of $\bar{\triangle}_{h_{\triangle,L}(v)}(v)$, there exists a point $p_i \in L$ such that $p_i \in F_i$ and p_i is not in any of F_j for $j \neq i$.*

Remark that this shows that every point in $\mathrm{Crit}(L)$ is a vertex of the Voronoi diagram $\mathrm{Vor}_{\triangle}(L)$.

Proof. We first prove that for every $v \in \mathrm{Crit}(L)$, there exist $(n+1)$ different points $p_i \in L$, $i = 0,\ldots,n$, such that the corresponding Voronoi cells $V_{\triangle}(p_i)$ shares v, i.e., such that $v \in V_{\triangle}(p_i)$ for $i \in \{\, 0,\ldots,n \,\}$. By Lemma 1.4.15, we know that there exists a point $x \in \mathrm{Ext}^c(L)$ such that $\pi_0(x) = v$. We will prove the following: there exist $(n+1)$ different points $p_i \in L$, $i = 0,\ldots,n$ such that $x \in H^-_{p_i}$ for all $i \in \{\, 0,\ldots,n \,\}$. Once this has been proved, we will be done. Indeed by Lemma 1.4.14, we know that that every Voronoi cell $V_\triangle(p)$, for $p \in L$, is of the form $\pi_0(H_p^-) \cap \partial\Sigma^c(L)$. So $v \in \pi_0(H_{p_i}^- \cap \partial\Sigma^c(L)) = V_\triangle(p_i)$ for each point p_i, and this is exactly what we wanted to prove.

To prove the second part, it will be enough to show that the points p_i have the desired property. Remark that we have $d_{\bar\triangle}(p_i, v) = d_\triangle(v, p_i) = h_{\triangle,L}(v)$, so $p_i \in \partial\bar\triangle_{h_{\triangle,L}(v)}(v)$ for all i. By the choice of p_i, we have $(p_i)_j > x_j$ for all $j \neq i$ and $(p_i)_i = x_i$. Since $v = \pi_0(x)$, we know that p_i is in the facet F_i of $\bar\triangle_{h_{\triangle,L}(v)}(v)$ defined by

$$F_i = \{\, u \in \bar\triangle_{h_{\triangle,L}(v)}(v) \;\mid\; u_i = v_i - h_{\triangle,L}(v) \text{ and } u_j \geq v_j - h_{\triangle,L}(v) \,\}.$$

(Remark that $d_{\bar\triangle}(x, v) = |\oplus_j (x_j - v_j)|$ so this is a facet of $\bar\triangle_{h_{\triangle,L}(v)}(v)$.) And p_i is not in any of the other facets F_j (since $(p_i)_j > v_j - h_{\triangle,L}(v)$ for $j \neq i$). So the proof of one direction is now complete. To prove the other direction, let v be a point such that each of the $n+1$ facets F_i of $\bar\triangle_{h_{\triangle,L}(v)}(v)$ has a point $p_i \in L$ and p_i is not in any of the other facets F_j for $j \neq i$. We show that v is critical, i.e., v is a local maxima of $h_{\triangle,L}$. It will be enough to show that for any non-zero vector $d \in H_0$ of sufficiently small norm, there exists one of the points p_i such that $d_\triangle(v+d, p_i) < h_{\triangle,L}(v) = d_\triangle(v, p_i)$. For all j, by the characterization of the facet F_j (see above) and by $p_j \notin F_k$ for all $k \neq i$, we have $d_\triangle(v+d, p_j) = d_{\bar\triangle}(p_j, v+d) = |\bigoplus_k (p_j)_k - v_k - d_k| = d_j + v_j - (p_j)_j = h_{\triangle,L}(v) + d_j$ if all d_k's are sufficiently small (namely if for all k, $|d_k| \leq \epsilon$ where $\epsilon > 0$ is chosen so that $2\epsilon < \min_{j,k:k\neq j}\big[(p_j)_k - v_k + h_{\triangle,L}(v)\big]$). As $d \in H_0$ and $d \neq 0$, there exists i such that $d_i < 0$. It follows that $h_{\triangle,L}(d+v) \leq d_\triangle(v+d, p_i) < h_{\triangle,L}(v)$. And this shows that v is a local maximum of $h_{\triangle,L}$. The proof of the lemma is now complete. □

1.4.5 Proof of Lemma 1.3.7

We end this section by providing the promised short proof of Lemma 1.3.7, which claims that the degree function is bounded below in the region $\Sigma^c(L)$.

In Section 1.4.4 we obtained the following explicit formula for $f_p(y)$:

$$\forall y \in H_0,\; f_p(y) \;=\; y - d_\triangle(y,p)(1,\ldots,1).$$

We infer that

$$\forall\, y \in V_\triangle(p): \quad f_p(y) \;=\; y - h_{\triangle,L}(y).(1,\ldots,1). \tag{1.5}$$

By Lemma 1.4.12, we have $\partial\Sigma^c(L) = \mathrm{Gr}(h_{\triangle,L})$. It follows from Equation 1.5 that

$$\partial\Sigma^c(L) \;=\; \{\, f_p(y) \mid y \in V_\triangle(p) \text{ and } p \in L\}.$$

We now observe that:

$$\forall\, y \in H_0: \; \deg(f_p(y)) \;=\; \deg(y) - (n+1)d_\triangle(y,p) \;=\; -\,(n+1)d_\triangle(y,p).$$

This shows that $\deg(f_p(y))$ depends only on the simplicial distance $d_\triangle$ between y and p. By translation invariance of the simplicial distance function (Lemma 1.4.2), translation invariance of the Voronoi cells (Lemma 5.1.7), and the above observations, we obtain

$$\begin{aligned}\inf(\deg(\Sigma^c(L))) &= \inf_{y\in V_\triangle(p)}\{\,-(n+1)d_\triangle(y,p)\,\}\\ &= \inf_{y\in V_\triangle(O)}\{\,-(n+1)d_\triangle(y,O)\,\}\\ &= -(n+1)\sup_{y\in V_\triangle(O)}\{\,d_\triangle(y,O)\,\}.\end{aligned}$$

By Lemma 1.4.7, we know that $V_\triangle(O)$ is compact. Also the function $d_\triangle(O,y)$ is continuous on y. Hence $\sup_{y\in V_\triangle(O)}\{d_\triangle(y,O)\}\}$ is finite and the lemma follows.

1.5 Uniform Reflection Invariant Sublattices

Consider a full dimensional sublattice L of A_n and its Voronoi diagram $\mathrm{Vor}_\triangle(L)$ under the simplicial distance function. From the previous sections, we know that the points of $\mathrm{Crit}(L)$ are vertices of $\mathrm{Vor}_\triangle(L)$. We know that $V_\triangle(O)$ is a compact star-shaped set with O as a kernel, and that the other cells are all translations of $V_\triangle(O)$ by points in L. Consider now the subset $\mathrm{Crit}V_\triangle(O)$ of vertices of $V_\triangle(O)$ which are in $\mathrm{Crit}(L)$. The sublattices of A_n of interest for us should have the following symmetry property:

Definition 1.5.1 (Reflection Invariance)**.** *A sublattice $L\subseteq A_n$ is called reflection invariant if $-\mathrm{Crit}(L)$ is a translate of $\mathrm{Crit}(L)$, i.e., if there exists $t\in\mathbb{R}^{n+1}$ such that $-\mathrm{Crit}(L)=\mathrm{Crit}(L)+t$. Furthermore, L is called strongly reflection invariant if the same property holds for $\mathrm{Crit}V_\triangle(O)$, i.e., if there exists $t\in\mathbb{R}^{n+1}$ such that $-\mathrm{Crit}V_\triangle(O)=\mathrm{Crit}V_\triangle(O)+t$.*

By translation invariance, we know that every strongly reflection invariant sublattice of A_n is indeed reflection invariant. Also, note that the vector t in the definition of reflection invariance lattices above is not uniquely defined: by translation invariance, if t' is linearly equivalent to t, t' also satisfies the property given in the definition.

Reflection Invariance and Involution of $\mathrm{Ext}(L)$. Let L be a reflection invariant sublattice and $t\in\mathbb{R}^{n+1}$ be a point such that $-\mathrm{Crit}(L)=\mathrm{Crit}(L)+t$. This means that for any $c\in\mathrm{Crit}(L)$ there exists a unique $\bar{c}\in\mathrm{Crit}(L)$ such that $c+\bar{c}=-t$. By Lemma 1.4.15 and Corollary 1.4.16, for every point c in $\mathrm{Crit}(L)$, there exists a point ν in $\mathrm{Ext}(L)$ such that $c=\pi_0(\nu)$. Thus, for every point ν in $\mathrm{Ext}(L)$, there exists a point $\bar{\nu}$ in $\mathrm{Ext}(L)$ such that $\pi_0(\nu+\bar{\nu})=-t$. This allows to define an involution $\phi(=\phi_t):\mathrm{Ext}(L)\to\mathrm{Ext}(L)$:

$$\text{For any point } \nu\in\mathrm{Ext}(L),\ \phi(\nu):=\bar{\nu}.$$

Note that ϕ is well defined. Indeed, if there exist two different points $\bar{\nu_1}$ and $\bar{\nu_2}$ such that $\pi_0(\nu+\bar{\nu_i})=-t$ for $i=1,2$, then $\pi_0(\bar{\nu_1})=\pi_0(\bar{\nu_2})$ and this would imply that $\bar{\nu_1}>\bar{\nu_2}$ or $\bar{\nu_2}>\bar{\nu_1}$ which contradicts the hypothesis that $\bar{\nu_1},\bar{\nu_2}\in\mathrm{Ext}(L)$. A similar argument shows that ϕ is a bijection on $\mathrm{Ext}(L)$ and is an involution.

1.5.1 A Riemann-Roch Inequality for Reflection Invariant Sub-Lattices

In this subsection, we prove a Riemann-Roch inequality for reflection invariant sublattices of A_n. We refer to Section 1.2.3 for the definition of g_{min} and g_{max}.

Let L be a reflection invariant sublattice of A_n. We have to show the existence of a canonical point $K \in \mathbb{Z}^{n+1}$ such that for every point $D \in \mathbb{Z}^{n+1}$, we have

$$3g_{min} - 2g_{max} - 1 \;\leq\; r(K-D) - r(D) + \deg(D) \;\leq\; g_{max} - 1\,. \tag{1.6}$$

K is defined up to linear equivalence (which is manifested in the choice of t in the definition of reflection invariance).

Construction of a Canonical Point K.

We define the canonical point K as follows: Let $\nu_0 \in \mathrm{Ext}(L)$ be an extremal point such that $\nu_0 + \phi(\nu_0)$ has the maximum degree, i.e., $\nu_0 = \mathrm{argmax}\{\deg(\nu + \phi(\nu)) \,|\, \nu \in \mathrm{Ext}(L)\}$. The map ϕ is the involution defined above. Define $K := -\nu_0 - \phi(\nu_0)$.

Theorem 1.5.2. (Weak Riemann-Roch) *Let L be a reflection invariant sublattice of A_n of rank n. There exists a point $K \in \mathbb{Z}^{n+1}$, called* canonical point*, such that for every point $D \in \mathbb{Z}^{n+1}$, we have*

$$3g_{min} - 2g_{max} - 1 \;\leq\; r(K-D) - r(D) + \deg(D) \;\leq\; g_{max} - 1\,.$$

Proof. We first observe that K is well-defined and for any point ν in $\mathrm{Ext}(L)$, $\nu + \bar{\nu} \leq -K$. This is true because all the points $\nu + \bar{\nu}$ are on the line $-t + \alpha(1,\ldots,1)$, $\alpha \in \mathbb{R}$, and K is chosen in such a way to ensure that $-K$ has the maximum degree among the points of that line. We infer that for any point $\nu \in \mathrm{Ext}(L)$, there exists an effective point E_ν such that $\nu + \bar{\nu} = -K - E_\nu$. Using this, we first derive an upper bound on the quantity $\deg^+(K - D + \bar{\nu}) - \deg^+(\nu + D)$ as follows:

$$\begin{aligned}
\deg^+(K - D + \bar{\nu}) - \deg^+(\nu + D) &= \deg^+(-\nu - \bar{\nu} - E_\nu - D + \bar{\nu}) - \deg^+(\nu + D) && (1.7)\\
&= \deg^+(-\nu - E_\nu - D) - \deg^+(\nu + D) && (1.8)\\
&\leq \deg^+(-\nu - D) - \deg^+(\nu + D) && (1.9)\\
&= \deg(-\nu - D) \;=\; -\deg(\nu) - \deg(D) && (1.10)\\
&\leq g_{max} - \deg(D) - 1. && (1.11)
\end{aligned}$$

To obtain Inequality (1.8), we use the fact that if $E \geq 0$ then $\deg^+(D - E) \leq \deg^+(D)$. Also remark that Inequality (1.11) is a simple consequence of the definition of g_{max}.

Now, we obtain a lower bound on the quantity $\deg^+(K - D + \bar{\nu}) - \deg^+(\nu + D)$. In order to do so, we first obtain an upper bound on the degree of E_ν, for the effective point E_ν such that $\nu + \bar{\nu} = -K - E_\nu$. To do so, we note that by the definition of K and by the definition of g_{min}, we have $\deg(K) = \min(\deg(-\nu - \bar{\nu})) \geq 2g_{min} - 2$. Also observe that by the definition of g_{max}, we have $\deg(-\nu - \bar{\nu}) \leq 2g_{max} - 2$. It follows that

$$\deg(E_v) = -\deg(K) + \deg(-\nu - \bar{\nu}) \leq 2(g_{max} - g_{min}).$$

We proceed as follows

$$\begin{aligned}
\deg^+(K - D + \bar{\nu}) - \deg^+(\nu + D) &= \deg^+(-\nu - E_\nu - D) - \deg^+(\nu + D) \\
&\geq \deg^+(-\nu - D) - \deg(E_\nu) - \deg^+(\nu + D) \\
&\geq 2(g_{min} - g_{max}) + \deg^+(-\nu - D) - \deg^+(\nu + D) \\
&\geq 2(g_{min} - g_{max}) - \deg(\nu + D) \\
&= 2(g_{min} - g_{max}) - \deg(\nu) - \deg(D) \\
&\geq 3g_{min} - 2g_{max} - \deg(D) - 1.
\end{aligned}$$

The last inequality follows from the definition of g_{min}. Now since the map $\phi(\nu) = \bar{\nu}$ is a bijection from $\mathrm{Ext}(L)$ onto itself, we can easily see that

$$\begin{aligned}
3g_{min} - 2g_{max} - \deg(D) - 1 &\leq \min_{\nu \in \mathrm{Ext}(L)} \deg^+(K + \bar{\nu} - D) - \min_{\nu \in \mathrm{Ext}(L)} \deg^+(\nu + D) \\
&\leq g_{max} - \deg(D) - 1.
\end{aligned}$$

By Lemma 1.2.9 and the fact ϕ is a bijection, we know that:

$$\begin{aligned}
r(D) &= \min\nolimits_{\nu \in \mathrm{Ext}(L)} \deg^+(\nu + D) - 1, \\
r(K - D) &= \min\nolimits_{\bar{\nu} \in \mathrm{Ext}(L)} \deg^+(K - D + \bar{\nu}) - 1.
\end{aligned}$$

Finally we infer that $3g_{min} - 2g_{\max} - \deg(D) - 1 \leq r(K - D) - r(D) \leq g_{max} - \deg(D) - 1$, and the Riemann-Roch Inequality follows. $\square$

Remark 1.5.3. As the above proof shows, we indeed obtain a slightly stronger inequality

$$g_{min} - \deg(D) - 1 - \max_{\nu \in \mathrm{Ext}(L)} \deg(E_\nu) \ \leq \ r(K - D) - r(D).$$

In particular if $E_\nu = 0$ for all $\nu \in \mathrm{Ext}(L)$ (see Section 2.2 for example regular digraphs), we have:

$$g_{min} - \deg(D) - 1 \ \leq \ r(K - D) - r(D) \leq g_{max} - \deg(D) - 1.$$

We remark that the proof technique used above is quite similar to the one used by Baker and Norine [12].

Remark 1.5.4. From Lemma 1.2.9, it is easy to obtain the inequality $\deg(D) - r(D) \leq g_{max}$, for all sublattices L of A_n and all $D \in \mathbb{Z}^{n+1}$. This inequality is usually referred to as Riemann's inequality. Note that the Riemann-Roch inequality (1.6) is more sensitive to (and contains more information about) the extent of "un-evenness" of the extremal points, while the above trivial inequality does not provide any such information.

1.5.2 Riemann-Roch Theorem for Uniform Reflection Invariant Lattices

Recall that a lattice L is called uniform if $g_{max} = g_{min}$, i.e., every point in $\mathrm{Ext}(L)$ has the same degree. By Corollary 1.4.16 and the definition of $h_\triangle$, this is equivalent to saying that the set of critical values of $h_{\triangle,L}$ is a singleton. We call $g = g_{max} = g_{min}$ the *genus* of the lattice.

The following is a direct consequence of Theorem 1.5.2. However we give it as a separate theorem.

Theorem 1.5.5. *Every uniform reflection invariant sublattice $L \subseteq A_n$ of dimension n has the Riemann-Roch property.*

Proof. Let $D \in \mathbb{Z}^{n+1}$. If L is a reflection invariant lattice, we can apply Theorem 1.5.2 to obtain $3g_{min} - 2g_{max} - 1 \leq r(K-D) - r(D) + \deg(D) \leq g_{max} - 1$, where K is the canonical point defined as in the proof of Theorem 1.5.2. Since L is uniform we have $g_{max} = g_{min} = g$ and we obtain $r(K-D) - r(D) + \deg(D) = g - 1$. It remains to show that $\deg(K) = 2g-2$. But, we know from the construction of K that $K = -(\nu+\bar{\nu})$ for a point $\nu \in \mathrm{Ext}(L)$. Since L is uniform, we infer that $\deg(K) = -\deg(\nu) - \deg(\bar{\nu}) = 2g-2$ (and also that $K = -\nu - \bar{\nu}$, $\forall\, \nu \in \mathrm{Ext}(L)$). $\square$

We say that a sublattice L of A_n has a Riemann-Roch formula if there exists an integer m and an integral point K_m, or simply K, of degree $2m-2$ (a canonical point) such that for every integral point D, we have:

$$r(D) - r(K-D) = \deg(D) - (m-1).$$

The following result shows the amount of geometric information one can obtain from the Riemann-Roch Property.

Theorem 1.5.6. *A sublattice L has a Riemann-Roch formula if and only if it is uniform and reflection invariant. Moreover, for a uniform and reflection invariant lattice $m = g$ (the genus of the lattice).*

The rest of this section is devoted to the proof of this theorem. One direction is already shown, we prove the other direction.

We first prove that

Claim 1.5.7. *If L has a Riemann-Roch formula, then $m = g_{max}$.*

Proof. The Riemann-Roch formula for a point D with $\deg(D) > 2m-2$ implies that $\deg(D) - r(D) = m$. On the other hand, using the formula for rank obtained in Lemma 1.2.9, we also know that $\deg(D) > 2g_{max} - 2$ then $\deg(D) - r(D) \leq g_{max}$. This for D with $\deg(D) \geq 2\max\{m, g_{max}\} - 2$ shows that $m \leq g_{max}$. By the Riemann-Roch formula, we have $r(D) \geq 0$ for any D with $\deg(D) \geq m$. Let $D = -\nu_{max}$, where ν_{max} is an extremal point of minimal degree. Remark that we have $r(D) = -1$. This shows that $m \geq g_{max}$. And we infer that $m = g_{max}$. $\square$

We now prove that

Claim 1.5.8. *If L has a Riemann-Roch formula, then L is uniform and $m = g$.*

Proof. Let N be the set of points of $\Sigma(L)$ of degree $-g_{max}+1$. We note that every point in N is extremal, i.e., $N \subset \mathrm{Ext}(L)$. To prove the uniformity, we should prove that $N = \mathrm{Ext}(L)$. We claim that $\Sigma(L) = \cup_{\nu \in N} H_\nu^+$, and this in turn implies that $N = \mathrm{Ext}(L)$. Indeed, if the claim holds, then every extremal point $\nu \in \mathrm{Ext}(L)$ should dominate a point u in N, and so $u = \nu$, meaning that $N = \mathrm{Ext}(L)$.

To prove the claim, we proceed as follows. Let $-D$ be a point in $\Sigma(L)$. We know that $r(D) = -1$. We should prove the existence of a point ν in N such that $\nu \leq -D$. We now claim that there exists $E \geq 0$ with $\deg(E) = g_{max} - 1 - \deg(D)$ and $r(D+E) = -1$. Assume the contrary, then for every point $E \geq O$ such that $deg(E) = g_{max} - 1 - deg(D)$, we have $r(D+E) \geq 0$. By the Riemann-Roch formula on the divisor $D+E$, we have $r(K-D-E) \geq 0$ and hence $r(K-D) \geq g_{max} - deg(D) - 1$. Now, using the Riemann-Roch formula on the divisor $K-D$, we have: $r(D) \geq 0$. A contradiction. The point $-D-E$ has degree $-g_{max}+1$ and so is in N. In addition $-D-E \leq -D$. And this is what we wanted to prove. The proof of the uniformity is now complete. □

To finish the proof of the theorem, it remains to show that

Claim 1.5.9. *If a uniform sublattice L of A_n of full rank has a Riemann-Roch formula, then it is reflection invariant.*

Proof. Consider a uniform lattice satisfying the Riemann-Roch property. By Lemma 1.2.2, we know that for a point ν in $\mathrm{Ext}(L)$, $r(-\nu) = -1$. Now, if we evaluate the Riemann-Roch formula for $D = -\nu$, we get $r(-\nu) = r(K+\nu)$. Hence, we have $r(-\nu) = r(K+\nu) = -1$. Again by Lemma 1.2.2, this implies that $-K-\nu$ is a point in $\Sigma(L)$. By the Riemann-Roch property and Claim 1.5.8 above, $\deg(K) = 2g-2$. Since L is uniform and $\nu \in \mathrm{Ext}(L)$, we have $\deg(\nu) = g-1$. We infer that $\deg(K+\nu) = g-1$, and it follows that $-K-\nu$ is an extremal point of $\Sigma(L)$. We now define $\bar{\nu} = -K-\nu$. Clearly, the map $\nu \to \bar{\nu}$ is a bijection from $\mathrm{Ext}(L)$ onto itself. Let $t = \pi_0(-K)$. We obtain $t = \pi_0(-\nu-\bar{\nu}) = -\pi_0(\nu) - \pi_0(\bar{\nu})$ for all $\nu \in \mathrm{Ext}(L)$. By Corollary 1.4.16, we have $\mathrm{Crit}(L) = \pi_0(\mathrm{Ext}(L))$ and hence $t = -c-\bar{c}$ for every c in $\mathrm{Crit}(L)$. This implies that $-\bar{c} = t+c$. To finish the proof, observe that $\bar{c} \to c$ is a bijection from $\mathrm{Crit}(L)$ onto itself, and so we have $-\mathrm{Crit}(L) = \mathrm{Crit}(L) + t$. □

The proof of Theorem 1.5.6 is now complete.

Chapter 2

Examples

In this chapter, we study the machinery that we presented in the previous chapter through various classes of examples. We start with the class of examples that originally motivated our work: Laplacian lattices of undirected graphs.

2.1 Lattices Generated by the Laplacian matrix of Connected Graphs

Probably the most interesting examples of the sublattices of A_n are generated by Laplacian of connected multi-graphs (and more generally directed multi-graphs) on $n+1$ vertices. In this subsection, we provide a geometric study of these sublattices. We prove the following result:

Theorem 2.1.1. *For any connected graph G, the sublattice L_G of A_n generated by the Laplacian of G is strongly reflection invariant and uniform.*

Theorem 2.1.1 will be a direct consequence of Theorem 2.1.9 below. Combining this theorem with Theorem 1.5.5 gives the main result of [12].

Corollary 2.1.2. (Theorem 1.12 in [12]) *For any undirected connected graph G on $n+1$ vertices and with m edges, the Laplacian lattice L_G has the Riemann-Roch property. In addition, we have $g_{\max} = g_{\min} = m - n$ and the canonical point K is given by $(\delta_0 - 2, \delta_1 - 2, \ldots, \delta_n - 2)$ of $\mathbb{Z}^{n+1}$ where δ_i's are the degrees of the vertices of G.*

Remark 2.1.3. Using reduced divisors, and the results of [12], it is probably quite straightforward to obtain a proof of Theorem 2.1.1. (This is not surprising since, as we pointed out in the previous section, a lattice with a Riemann-Roch formula has to be uniform and reflection invariant.) The proof we will present for Theorem 2.1.1 indeed gives more than what is the content of this theorem. We give a complete description of the Voronoi-diagram and its dual Delaunay triangulation (we will give a precise definition in Chapter 3). And we do not use reduced divisors, which is the main tool used in the previous proofs of the Riemann-Roch theorem. As we will see, the form of the canonical divisor for a given graph (and the genus) as defined in [12] comes naturally out of this explicit description.

Let G be a connected graph on $n+1$ vertices $v_0, v_1, \ldots, v_n$ and m edges. Let L_G, or simply L if there is no risk of confusion, be the Laplacian sublattice of A_n. We summarise the main properties of the lattice L_G and the matrix Q. L_G is a rank n sublattice of A_n with $\{b_0, \ldots, b_{n-1}\}$ as a basis such that the $(n+1)\times(n+1)$ matrix Q has $\{b_0, \ldots, b_{n-1}\}$ as the first n rows and $b_n = -\sum_{i=0}^{n-1} b_i$ as the last row. In addition, the matrix

$$Q = \begin{bmatrix} \delta_0 & -b_{01} & -b_{02} \ldots & -b_{0n} \\ -b_{10} & \delta_1 & -b_{12} \ldots & -b_{1n} \\ \vdots & \vdots & \ddots & \\ -b_{n0} & -b_{n1} & -b_{n2} \ldots & \delta_n \end{bmatrix} \tag{2.1}$$

has the following properties:

(C_1) b_{ij}'s are integers, $b_{ij} \geq 0$ for all $0 \leq i \neq j \leq n$ and $b_{ij} = b_{ji}$, $\forall i \neq j$.

(C_2) $\delta_i = \sum_{j=1, j\neq i}^{n} b_{ij} = \sum_{j=1, j\neq i}^{n} b_{ji}$ (and is the degree of the i-th vertex).

We denote by B the basis $\{b_0, \ldots, b_{n-1}\}$ of L_G.

2.1.1 Voronoi Diagram $\mathrm{Vor}_\triangle(L_G)$ and the Riemann-Roch Theorem for Graphs

We first provide a decomposition of H_0 into simplices with vertices in L such that the vertices of each simplex forms an affine basis of L_G. Recall that a subset of lattice points $X \subset L$ of size $n+1$ is called an affine basis of L, if for $v \in X$, the set of vectors $u - v$, $u \in X$ and $u \neq v$, forms a basis of L. In other words, if the simplex defined by X is minimal (which means it is full-dimensional and has minimum volume among all the (full-dimensional) simplices whose vertices lie in L). The whole decomposition is derived from the symmetries of the affine basis B, and describes in a very nice way the Voronoi decomposition $\mathrm{Vor}_\triangle(L_G)$. What follows could be considered as an explicit construction of the "*Delaunay dual*", $\mathrm{Del}_\triangle(L_G)$ of $\mathrm{Vor}_\triangle(L_G)$.

We consider the family of total orders on the set $\{0, 1, \ldots, n\}$. A total order $<_\pi$ on $\{0, 1, \ldots, n\}$ gives rise to an element π of the symmetric group S_{n+1}, defined in such a way that $\pi(0) <_\pi \pi(1) <_\pi \cdots <_\pi \pi(n-1) <_\pi \pi(n)$. It is clear that the set of all total orders on $\{0, \ldots, n\}$ is in bijection with the elements of S_{n+1}. In addition, the total orders which have n as the maximum element are in bijection with the subgroup $S_n \subset S_{n+1}$ consisting of all the permutations which fix n, i.e., $\pi(n) = n$. In the following when we talk about a permutation in S_n, we mean a permutation of S_{n+1} which fixes n. For $\pi \in S_n$, we denote by $\bar{\pi}$ the *opposite permutation* to π defined as follows: we set $\bar{\pi}(n) = n$ and $\bar{\pi}(i) = \pi(n-1-i)$ for all $i = 0, \ldots, n-1$. In other words, for all $i = 0, \ldots, n-1$, $i <_\pi j$ if and only if $j <_{\bar{\pi}} i$, and $j \leq_{\bar{\pi}} n$ for all j. Let C_{n+1} denotes the group of cyclic permutations of $\{0, \ldots, n\}$, i.e., $C_{n+1} =< \sigma >$ where σ is the element of

S_{n+1} defined by $\sigma(i) = i+1$ for $0 \le i \le n-1$ and $\sigma(n) = 0$. It is easy to check that $S_{n+1} = S_n C_{n+1}$.

Let $<_\pi$ be a total order such that $\pi \in S_n$, i.e., $\pi \in S_{n+1}$ and $\pi(n) = n$. We first define a set of vectors $B^\pi = \{\, b_0^\pi, \dots, b_n^\pi \,\}$ as follows:

$$\forall\, i \in \{\, 0, \dots, n \,\}, \quad b_i^\pi := \sum_{j \le_\pi i} b_j \, .$$

In particular, note that $b_n^\pi = b_{\pi(n)}^\pi := \sum_{j \le_\pi \pi(n)} b_j = \sum_{j=0}^{n} b_j = 0$.

Lemma 2.1.4. *For any total order $<_\pi$ with n as maximum, or equivalently for any π in S_n, the set $B^\pi = \{\, b_0^\pi, \dots, b_n^\pi \,\}$ forms an affine basis of L_G.*

Proof. We first observe that the matrix of $\{b_{\pi(0)}^\pi, \dots, b_{\pi(n-1)}^\pi\}$ in the base B is upper triangular with diagonals equal to 1. It follows that the set $\{b_{\pi(0)}^\pi, \dots, b_{\pi(n-1)}^\pi\}$ is an affine basis of L. As $b_{\pi(n)}^\pi = 0$, it follows that B^π is a basis. $\square$

We denote by $\triangle^\pi$ the simplex defined by B^π. In other words, $\triangle^\pi := \mathrm{Conv}(B^\pi)$, the convex-hull of B^π. Consider the fundamental parallelotope $F(B)$ defined by the basis B of L_G. Note that $F(B)$ is the convex-hull of all the vectors b_i^π for $\pi \in S_n$ and $i \in \{0, \dots, n\}$. We next show that the set of simplices $\{\triangle^\pi\}_{\pi \in S_n}$ provides a simplicial decomposition (i.e., a triangulation) of $F(B)$. But before we need the following simple lemma attributed to Freudenthal:

Lemma 2.1.5. *Let $\square_n = \{\, (x_0 \dots, x_{n-1}) \,|\, 0 \le x_i \le 1\}$ be the unit hypercube in $\mathbb{R}^n$. For a permutation $\pi \in S_n$, let $\bar{\triangle}_n^\pi = \{\, x = (x_0, \dots, x_{n-1}) \in \mathbb{R}^n \mid 0 \le x_{\pi(n-1)} \le x_{\pi(n-2)} \le \dots \le x_{\pi(0)} \le 1\}$. The set of simplices $\{\, \bar{\triangle}_n^\pi \,\}_{\pi \in S_n}$ is a simplicial decomposition of $\square_n$.*

We have

Lemma 2.1.6. *Let G be a connected graph and $L \subset A_n$ be the corresponding Laplacian lattice. The set of simplices $\{\triangle^\pi\}_{\pi \in S_n}$ is a simplicial decomposition of $F(B)$.*

Proof. Since B is a basis of the rank n lattice L_G, which is contained in H_0, it is also a basis of H_0. By definition, $F(B)$ is the unit cube with respect to the basis B. By Lemma 2.1.5, the family of simplices $\{\bar{\triangle}^\pi\}_{\pi \in S_n}$ is a simplicial decomposition of $F(B)$, where $\bar{\triangle}^\pi = \{\, x = x_0 b_0 + \dots + x_{n-1} b_{n-1}) \in H_0 \,|\, 0 \le x_{\pi(n-1)} \le x_{\pi(n-2)} \le \dots \le x_{\pi(0)} \le 1\}$ and the vectors are written in the B-basis. Now recall that the vertices of $\triangle^\pi$ are given by the points b_j^π. Recall also that $\forall\, i \in \{\, 0, \dots, n \,\}$, $b_i^\pi := \sum_{j \le_\pi i} b_j$, and that $b_n^\pi = 0$. A simple calculation shows that $\triangle^\pi$ coincides with the simplex $\bar{\triangle}^\pi$ above, and the proof follows. $\square$

A combination of this lemma with the simple fact that $F(B) + L_G$ is a tiling of H_0 gives us:

Corollary 2.1.7. *The set of simplices $\{\, \triangle^\pi + p \,|\, \pi \in S_n,\, p \in L_G \,\}$ forms a triangulation of H_0.*

In the simplicial decomposition $\{\triangle^\pi + p \,|\, \pi \in S_n \,\&\, p \in L\}$ of H_0, consider the set $\mathcal{S}im_O$ consisting of all the simplices that contain the origin O as a vertex. We have

Lemma 2.1.8. *A simplex is in $\mathcal{S}im_O$ if and only if it is spanned by B^π for some π in S_{n+1}. (Remark that we do not assume that $\pi(n) = n$.)*

Proof. By Corollary 2.1.7, we know that every simplex in $\mathcal{S}im_O$ is of the form: $\triangle^{\pi_0} + q$ for some π_0 in S_n and q in L. Recall that the element π of S_n is regarded as an element of S_{n+1}, with the property that $\pi(n) = n$. Since, the vertex set of $\triangle^{\pi_0}$ is $V(\triangle^{\pi_0}) = \{b_0^{\pi_0}, \dots, b_{n-1}^{\pi_0}, O\}$, we should have $q = -b_i^{\pi_0}$ for some $0 \le i \le n-1$. Let $0 \le j \le n$ be such that $\pi_0(j) = i$. A straightforward calculation shows that $V(\triangle^{\pi_0}) - b_i^{\pi_0} = V(\triangle^\pi)$, where $\pi = \pi_0\sigma^j$ and σ is the cyclic permutation $(0, 1, 2, .., n) \to (1, \dots, n, 0)$. The lemma follows because every element $\pi \in S_{n+1}$ can be written uniquely in the form $\sigma^i \pi_0$ for some $\pi_0 \in S_n$ ($S_{n+1} = S_n C_{n+1}$). $\square$

Remark also that $|\mathcal{S}im_O| = |S_{n+1}| = (n+1)!$.

Our aim now will be to provide a complete description of the set $\mathrm{Ext}^c(L_G)$ of extremal points of $\Sigma^c(L)$ (and equivalently the set $\mathrm{Ext}(L_G) = \mathrm{Ext}^c(L_G) - (1, \dots, 1)$) in terms of this triangulation. Actually we obtain an explicit description of the set $\mathrm{Crit}V_\triangle(O)$. Before we proceed, let us introduce an extra notation. Let π be an element of the permutation group S_{n+1}. We do not suppose anymore that $\pi(n) = n$. We define the point $\nu^\pi \in \mathbb{Z}^{n+1}$ as the tropical sum of the points of B^π, i.e., $\nu^\pi := \bigoplus_{i=0}^n b_i^\pi$. (And recall that $b_i^\pi = \sum_{j \le_\pi i} b_j$.) We have the following theorem.

Theorem 2.1.9. *Let G be a connected graph and L_G be the Laplacian lattice of G.*

(i) The set of extremal points of $\Sigma^c(L_G)$ consists of all the points $\nu^\pi + p$ for $\pi \in S_{n+1}$ and $p \in L_G$, i.e., $\mathrm{Ext}^c(L_G) = \{\, \nu^\pi + p \,|\, \pi \in S_{n+1}$ and $p \in L_G \,\}$. As a consequence, we have $\mathrm{Ext}(L_G) = \{\, \nu^\pi + p + (1, \dots, 1) \,|\, \pi \in S_{n+1}$ and $p \in L_G \,\}$.

(ii) We have $\mathrm{Crit}V_\triangle(O) = \pi_0(\{\, \nu^\pi \,|\, \pi \in S_{n+1} \,\})$.

We verify that the set $\{\nu^\pi\}$ has the following properties (c.f. Theorem 2.1.1 below.)

(P1)- **Reflection Invariance.** For all $\pi \in S_{n+1}$, $\nu^\pi + \nu^{\bar\pi} = (-\delta_0, -\delta_1, \dots, -\delta_n)$ where $\bar\pi$ is the opposite permutation to π, and δ_i denotes the degree of the vertex v_i. Since $\mathrm{Crit}V_\triangle(O) = \pi_0(\{\, \nu^\pi \,|\, \pi \in S_{n+1} \,\})$, it follows that L_G is strongly reflection invariant. More precisely we have $\mathrm{Crit}V_\triangle(O) = -\mathrm{Crit}V_\triangle(O) + \pi_0((-\delta_0, \dots, -\delta_n))$. (Recall that π_0 is the projection function.)

(P2)- **Uniformity.** For all $\pi \in S_{n+1}$, $\deg(\nu^\pi) = -m$. In other words, the Laplacian lattice L_G is uniform.

The proof of the results of this section will be given in the next subsection. However, let us quickly show how to calculate g and K in the above corollary. The vertices ν^π all belong to Ext^c and have degree $-m$. It follows that the vertices of $\mathrm{Ext}(L) = \mathrm{Ext}^c + (1, \dots, 1)$ have all degree $-m + n + 1$, and so by the definition of genus, we obtain $g_{min} = g_{max} = m - n$. In particular g coincides with the graphical genus of G (which is the number of vertices minus the number of edges plus one). Since the points of $\mathrm{Ext}(L_G)$ are of the form $\nu^\pi + (1\ldots,1)$, and as we saw in the proofs of Theorem 1.5.2 and Theorem 1.5.5, we have $K = -(\nu^\pi + (1,\ldots,1)) - (\nu^{\bar\pi} + (1,\ldots,1)) = (\delta_0 - 2, \delta_1 - 2, \ldots, \delta_n - 2)$.

2.1.2 Proofs of Theorem 2.1.9 and Theorem 2.1.1

Observe that the point $\nu^\pi = \bigoplus_{i=0}^n b_i^\pi$ has the following explicit form:

$$\nu^\pi \;=\; (-\sum_{j<_\pi 0} b_{j0}, -\sum_{j<_\pi 1} b_{j1}, \ldots, -\sum_{j<_\pi n} b_{nj}). \tag{2.2}$$

It follows that

$$\begin{aligned}
\nu^\pi \;&=\; (-\delta_0 + \sum_{j>_\pi 0} b_{j0}, -\delta_1 + \sum_{j>_\pi 1} b_{j1}, \ldots, -\delta_n + \sum_{j>_\pi n} b_{nj}) \\
&=\; (-\delta_0, \ldots, -\delta_n) - (-\sum_{j<_{\bar\pi} 0} b_{j0}, -\sum_{j<_{\bar\pi} 1} b_{j1}, \ldots, -\sum_{j<_{\bar\pi} n} b_{nj}) \\
&=\; (-\delta_0, \ldots, -\delta_n) - \nu^{\bar\pi}.
\end{aligned}$$

And we infer that

Lemma 2.1.10. *For every $\pi \in S_{n+1}$, we have $\nu^\pi + \nu^{\bar\pi} = (-\delta_0, \ldots, -\delta_n)$.*

Second, we calculate the degree of the point ν^π. We compute

$$\deg(\nu^\pi) = -\sum_{i,j\,:\,j<_\pi i} b_{ij} = -m,$$

where m denotes the number of edges of G, or equivalently in terms of the matrix Q, $m = \frac{1}{2}\sum_i \delta_i = trace(Q)/2$. It follows that

Lemma 2.1.11. *All the points ν^π have the same degree.*

We now show that $\nu^\pi \in \Sigma^c(L)$ for every $\pi \in S_{n+1}$. Assume for the sake of contradiction that there exists a point $p \in L$ such that $p > \nu^\pi$. By the definition of the Laplacian lattice L, we know that there are integers $\alpha_0, \ldots, \alpha_n$ such that $p = \alpha_0 b_0 + \ldots, \alpha_n b_n$, and so we can write

$$p = (\sum_{j=0}^n (\alpha_0 - \alpha_j) b_{j0}, \ldots, \sum_{j=0}^n (\alpha_n - \alpha_j) b_{jn}).$$

for some $\alpha_i \in \mathbb{Z}$. Among the integer numbers α_i, consider the set of indices S_p consisting of the indices i for which α_i is minimum. Remark that as p is certainly non zero (since there is a coordinate of ν^π which is zero, we cannot have $0 > \nu^\pi$), we cannot have $S_p = \{0, \ldots, n\}$. Now in the set S_p consider the index k which is the minimum in the total order $<_\pi$. By construction of k, we have $\alpha_k - \alpha_j \leq -1$ for all $j <_\pi k$ and $\alpha_k - \alpha_j \leq 0$ for all $j \geq_\pi k$. It follows that p_k, the $k-$th coordinate of p, is bounded above by

$$p_k \;=\; \sum_{j=0}^{n} (\alpha_k - \alpha_j) b_{jk} \;\leq\; \sum_{j <_\pi k} -b_{jk} \;=\; \nu_k^\pi.$$

And this contradicts our assumption $p > \nu^\pi$.

Next, we need to show that ν^π is a local minimum of the degree function. We already know that $\deg(\nu^\pi) = -m$. We will prove that for every point $x \in \Sigma^c(L)$, we have $\deg(x) \geq -m$. By Lemma 1.4.12, it will be enough to prove that $h_{\triangle,L}(x) \leq \frac{m}{n+1}$ for every point $x \in L$. By the definition of the simplicial distance function $h_{\triangle,L}$, this is equivalent to proving that the simplex $x + \frac{m}{n+1}\triangle$ contains a lattice point $p \in L$, i.e.,

$$\forall x \in H_0,\;\; (x + \frac{m}{n+1}\triangle) \cap L \neq \emptyset. \tag{2.3}$$

Here we use the following perturbation trick to reduce the problem to the case when all the entries of Q are non-zero. We add a rational number $\epsilon = \frac{s}{t}$, $s,t \in \mathbb{N}$, to each b_{ij}, $i \neq j$, to obtain b_{ij}^ϵ. We also define δ_i^ϵ in such a way that $\sum_i b_{ij}^\epsilon = \delta_i^\epsilon$. Remark that $\sum_j \delta_j^\epsilon = \frac{tr(Q^\epsilon)}{2}$. The new matrix Q^ϵ is not integral anymore (but if we want to work with integral lattices, we can multiply every coordinate by a large integer t to obtain an integral matrix tQ^ϵ). If we know that our claim is true for all Laplacians with non-zero coordinates, then the function h associated to tQ^ϵ satisfies the property

$$h_{\triangle,L^{t,\epsilon}} \leq \frac{tr(tQ^\epsilon)}{2(n+1)}. \tag{2.4}$$

Where $L^{t,\epsilon}$ denotes the lattice generated by the matrix tQ^ϵ. Let L^ϵ be the (non necessarily integral) lattice generated by the matrix Q^ϵ. We have $t.h_{\triangle,L^\epsilon} = h_{\triangle,L^{t,\epsilon}}$. Equation 2.4 implies then

$$h_{\triangle,L^\epsilon} \leq \frac{tr(Q^\epsilon)}{2(n+1)} = \frac{m}{n+1} + \frac{n\epsilon}{2(n+1)}. \tag{2.5}$$

Using characterization of Equation 2.3, one can see that, varying ϵ, the above property for all sufficiently small rational $\epsilon > 0$ will imply that $h_{\triangle,L} \leq \frac{m}{n+1}$, and that is what we wanted to prove. Indeed one can easily show that the distance function $h_{\triangle,L_\epsilon}(p)$ is a continuous function in ϵ and p.

So at present, we have shown that we can assume that all the b_{ij}'s are strictly positive. This is the assumption we will make for a while. In this case, using the explicit calculation of ν^π, we have:

Lemma 2.1.12. *The point ν^π has the following properties:*
1. $\nu_i^\pi = b_{ii}^\pi$ for $0 \leq i \leq n$.
2. $\nu_j^\pi < b_{ij}^\pi$ for $i \neq j$ and $0 \leq i,j \leq n$.

As a corollary we obtain:

Corollary 2.1.13. *Let $\{e_0, \ldots, e_n\}$ be the standard orthonormal basis of $\mathbb{R}^n$, i.e., $e_0 = (1, 0, \ldots, 0), \ldots, e_n = (0, \ldots, 0, 1)$. Let e_j be a fixed vector. For every $\delta > 0$, $\nu^\pi - \delta e_j < b^\pi_j$.*

Furthermore, as a direct consequence of Corollary 2.1.13 we obtain:

Lemma 2.1.14. *For every non-zero vector w in H^+_O and for every $\delta > 0$, there exists a point p in L such that $\nu^\pi - \delta w < p$.*

It follows now easily that

Corollary 2.1.15. *The point ν^π is an extremal point of $\Sigma^c(L)$.*

Proof. Follows by combining Lemmas 2.1.14 and 1.3.6. □

We will now prove the following: *every extremal point of $\Sigma^c(L)$ can be written as the tropical sum of the vertices of a simplex of the form $\triangle^\pi + q$, for some $\pi \in S_n$ and some q in L.* Again we will first assume a stronger condition that $b_{ij} > 0$ for all i, j such that $i \neq j$ and $1 \leq i, j \leq n-1$. And then we do a limiting argument similar to the one we did above to obtain the general statement. Let $\pi \in S_n$ a fixed permutation. Using the assumption $b_{ij} > 0$ for $i \neq j$, we have the following property:

Lemma 2.1.16. *For any total ordering $<_\pi$, $\pi \in S_{n+1}$, we have $b^\pi_{ij} \neq 0$ for all $0 \leq i, j \leq n$ and $i \neq \pi(n)$. (Remark that $b^\pi_{\pi(n)} = 0$.) Here b^π_{ij} is the $j-$th coordinate of the vector b^π_i. In addition, if $b^\pi_{ij} > 0$ (resp. $b^\pi_{ij} < 0$), then $j \leq_\pi i$ (resp. $i <_\pi j <_\pi \pi(n)$).*

As we saw in Lemma 2.1.8, the set of simplices Δ^π, $\pi \in S_{n+1}$, coincides with $\mathcal{S}im_O$, the set of all simplices of the triangulation which are adjacent to O. The simplices of $\mathcal{S}im_O$ naturally define a fan $\mathcal{F}$, the maximal elements of which are the set of all cones $\mathcal{C}^\pi$ generated by Δ^π for $\pi \in S_{n+1}$. In other words if B^π denoted the affine basis $\{b^\pi_i\}^n_{i=0}$, the cone $\mathcal{C}^\pi$ is the cone generated by B^π. In particular every element of H_0 is in some $\mathcal{C}^\pi$ for some $\pi \in S_{n+1}$. We have

Lemma 2.1.17. *Let q be a point in L, and $q \neq b^\pi_i$ for all $\pi \in S_{n+1}$ and $0 \leq i \leq n$. Let $\mathcal{C}^\pi$ be a cone in $\mathcal{F}$ which contains q. There exists a vector b^π_i in B^π such that $p < b^\pi_i$ for every point p in $H^-_O \cap H^-_q$. In particular, no point in $H^-_O \cap H^-_q$ is contained in $\Sigma^c(L)$.*

Proof. Since q is a point in $L \cap \mathcal{C}^\pi$, there exists non-negative integers $\alpha_i \geq 0$, $0 \leq k \leq n-1$, such that we can write $q = \sum_{k=0}^{n-1} \alpha_k b^\pi_{\pi(k)}$. In addition, since $q \notin B^\pi$, we have $\sum_l \alpha_l \geq 2$. Let $j = \min\{\, k \mid \alpha_k \neq 0 \,\}$, i.e., the minimum index such that $\alpha_k \neq 0$, and let $i = \pi(j)$. We show that the point b^π_i satisfies the condition of the lemma. For this, it will be enough to prove that $b^\pi_i > O \oplus q$. Indeed $p \in H^-_O \cap H^-_q$ implies that $p \leq O \oplus q$, and so if $b^\pi_i > O \oplus q$, then we have $p < b^\pi_i$, which is the required claim.

We should prove that $b^\pi_{ik} > (O \oplus q)_k$ for all k. As $i = \pi(j) \neq \pi(n)$, by Lemma 2.1.16 we know that $b^\pi_{ik} \neq 0$ for all k. There are two cases: if $b^\pi_{ik} > 0$, then easily we have

$b^\pi_{ik} > 0 \geq (O \oplus q)_k$. If $b^\pi_{ik} < 0$, then by Lemma 2.1.16, we have $b^\pi_{lk} < 0$ for all $l \geq_\pi i$. By the choice of i, we have $\alpha_l = 0$ for all $l <_\pi i$. We infer that $b^\pi_{jk} > \sum_l \alpha_l b^\pi_{lk} = (O \oplus q)_k$, and the lemma follows. □

We obtain the following corollary: *the simplices of our simplicial decomposition form the dual of the Voronoi diagram.* More precisely

Corollary 2.1.18. *Let q be a point in L that is not a vertex of a simplex in $\mathcal{S}im_O$, i.e., $q \neq b^\pi_i$ for all $\pi \in S_{n+1}$ and $0 \leq i \leq n$. Then $V(O) \cap V(q) = \emptyset$. Hence, for every two points p and q in L, we have $V(p) \cap V(q) \neq \emptyset$ if and only if p and q are adjacent in the simplicial decomposition of H_0 defined by $\{\triangle^\pi + p \,|\, \pi \in S_n \,\&\, p \in L\}$ i.e., $V(p) \cap V(q) \neq \emptyset$ if and only if there exists $\pi \in S_{n+1}$ such that q is a vertex of $\Delta^\pi + p$.*

Proof. We prove the first statement by contradiction. So for the sake of a contradiction, assume the contrary and let $p \in V(O) \cap V(q)$. By definition, we have $h_{\triangle,L}(p) = d_\triangle(p, O) = d_\triangle(p, q) \leq d_\triangle(p, q')$ for all points $q' \in L$. By Lemma 1.4.14 this implies that the point $y = f_O(p) = f_q(p)$ is a point in $\partial\Sigma^c(L)$ (c.f. Section 1.4 for the definition of f_p). By the definition of f_p, the point y is in $H^-_O \cap H^-_q$. On the other hand, Lemma 2.1.17 implies that no point in $H^-(O) \cap H^-(q)$ can be contained in $\Sigma^c(L)$. We obtain a contradiction. To see the second part, by translation invariance we can assume $p = O$. And in this case, the results follows by observing that for $q \in \Delta^\pi$, $\pi_0(\nu^\pi) \in V(O) \cap V(q)$. □

We can now present the proof of Theorem 2.1.9 in the case where all the b_{ij}'s are strictly positive. It will be enough to prove that $\mathrm{Crit}V(O) = \pi_0(\{\nu^\pi \mid \pi \in S_{n+1}\})$. As $v_\pi \in H^-_O$ and we showed that v_π is in $\mathrm{Ext}^c(L)$, we have $\pi_0(\{\nu^\pi \,|\, \pi \in S_{n+1}\}) \subseteq \mathrm{Crit}V(O)$. We show now $\mathrm{Crit}V(O) \subseteq \pi_0(\{\nu^\pi \mid \pi \in S_{n+1}\})$. Let $v \in \mathrm{Crit}V(O)$ and x be the point in $\mathrm{Ext}^c(L)$ with $\pi_0(x) = v$. By Lemma 1.4.17, there exist points $p_0, \ldots, p_n \in L$ such that $v \in V(p_0) \cap \cdots \cap V(p_n)$. By Corollary 2.1.18, points $p_0, \ldots, p_n$ should be adjacent in the simplicial decomposition of H_0 defined by $\{\triangle^\pi + p \mid \pi \in S_n \,\&\, p \in L\}$. As v is also in $\mathrm{Vor}(O)$, it follows that one of the p_i is O, and so there exists $\pi \in S_{n+1}$ such that $v \in \cap_{p \in B^\pi} \mathrm{Vor}(p)$. By the proof of Lemma 1.4.17, we also have $x = \bigoplus_i p_i$. But $\bigoplus_{p \in B^\pi} p = \nu^\pi$. It follows that $x = \nu^\pi$. We infer that $v \in \pi_0(\{\nu^\pi \mid \pi \in S_{n+1}\})$ and the theorem follows.

The proof of Theorem 2.1.1 is a simple consequence of Lemma 2.1.10, and what we just proved, namely, $\mathrm{Ext}^c(L) = \{\nu^\pi + q \,|\, \pi \in S_n \,\&\, q \in L\}$ and $\mathrm{Crit}V(O) = \pi_0(\{\nu^\pi \,|\, \pi \in S_{n+1}\})$.

To prove the general case, it will be enough to show that $\mathrm{Crit}V(O) = \pi_0(\{\nu^\pi \mid \pi \in S_{n+1}\})$ still holds. Indeed the rest of the arguments remain unchanged.

We consider again the ϵ-perturbed Laplacian Q_ϵ and do a limiting argument similar to the one we did before. Let L_ϵ to be the lattice generated by Q_ϵ. By $\mathrm{Vor}(L_\epsilon)$ and $\mathrm{Crit}V_\epsilon(p)$, we denote the Voronoi diagram of L_ϵ under the distance function $d_\triangle$ and the Voronoi cell of a point $p \in L_\epsilon$. We also define B^π_ϵ, Δ^π_ϵ, and ν^π_ϵ similarly.

Theorem 2.1.9 in the case where all the coordinates are strictly positive implies that $\mathrm{Crit}V_\epsilon(O) = \pi_0(\{v^\pi_\epsilon \mid \pi \in S_{n+1}\})$. We can naturally define limits of the sets $\mathrm{Crit}V_\epsilon(O)$

as ϵ tends to zero as limits of the points $\pi_0(v_\epsilon^\pi)$. Indeed this limit exists and coincides with the set $\pi_0(\{\nu^\pi \mid \pi \in S_{n+1}\})$, as can be easily verified. We show now

Lemma 2.1.19. *We have* $\lim_{\epsilon\to 0} \mathrm{Crit}V_\epsilon(O) = \mathrm{Crit}V(O)$.

Remark 2.1.20. Unfortunately this is not true in general for non graphical lattices. However, we always have $\mathrm{Crit}V(O) \subseteq \lim_{\epsilon\to 0}\mathrm{Crit}V_\epsilon(O)$.

Proof of Lemma 2.1.19. By Corollary 2.1.15, we already know that every point of $\pi_0(\{\nu^\pi|\pi \in S_{n+1}\})$ is critical. So we should only prove that these are the only critical points, namely $\mathrm{Crit}V(O) \subseteq \lim_{\epsilon\to 0}\mathrm{Crit}V_\epsilon(O) = \pi_0(\{\nu^\pi \mid \pi \in S_{n+1}\})$. Let c be a critical point of L. By Lemma 1.4.17, we know that there exists a set of points $p_0, \ldots, p_n$ such for each i, the facet F_i of $\bar{\triangle}_{h_{\triangle,L}(c)}(c)$ contains p_i and none of the other points $p_j \neq p_i$. We will show the following: for all sufficiently small ϵ, there exists a point $c_\epsilon \in L_\epsilon$ and $h_\epsilon = h_{\triangle,L_\epsilon}(c_\epsilon) \in \mathbb{R}_+$ such that $\bar{\triangle}_{h_\epsilon}(c_\epsilon)$ has the same property for the lattice L_ϵ, namely, for each i, the facet $F_{\epsilon,i}$ of $\bar{\triangle}_{h_\epsilon}(c_\epsilon)$ contains a point $p_{\epsilon,i} \in L_\epsilon$ which is not in any other facet $F_{\epsilon,j}$ of $\bar{\triangle}_{h_\epsilon}(c_\epsilon)$, for $j \neq i$. In addition $\bar{\triangle}_{h_\epsilon}(c_\epsilon) \to \bar{\triangle}_{h_{\triangle,L}(c)}(c)$, and so $h_\epsilon \to h_{\triangle,L}(c)$ and $c_\epsilon \to c$ (h_ϵ and c_ϵ being the radius and the centre of these balls $\bar{\triangle}_{h_\epsilon}(c_\epsilon)$). As each of the point c_ϵ will be critical for L_ϵ, we conclude that $c \in \lim_{\epsilon\to 0}\mathrm{Crit}(L_\epsilon)$ which is easily seen to be enough for the proof of the lemma. To show this last statement, we argue as follows: for small enough ϵ, there exist points $q_{\epsilon,0}$ and $p_{\epsilon,1}, \ldots, p_{\epsilon,n} \in L_\epsilon$ such that $q_{\epsilon,0} \to p_0$ and for all $n \geq i \geq 1$, $p_{\epsilon,i} \to p_i$ when ϵ goes to zero. These points naturally define a ball for the metric $d_{\bar{\triangle}}$, i.e., a simplex of the form $\bar{\triangle}_{r_\epsilon}(\bar{c}_\epsilon)$. This is the bounded simplex defined by the set of hyperplanes $E_{i,\epsilon}$, where $E_{i,\epsilon}$ is the hyperplane parallel to the facet F_i of $\bar{\triangle}_{h_{\triangle,L}(c)}(c)$ which contains $p_{\epsilon,i}$ ($q_{\epsilon,0}$ for $i = 0$). We define the ball $\bar{\triangle}_{h_\epsilon}(c_\epsilon)$ as follows. For each ϵ, if the interior of $\bar{\triangle}_{r_\epsilon}(\bar{c}_\epsilon)$ does not contain any other lattice point (a point of L_ϵ), we let $\bar{\triangle}_{h_\epsilon}(c_\epsilon) := \bar{\triangle}_{r_\epsilon}(\bar{c}_\epsilon)$. If the interior of $\bar{\triangle}_{r_\epsilon}(\bar{c}_\epsilon)$ contains another point of L_ϵ, let $p_{\epsilon,0}$ be the furthest point from the hyperplane $E_{0,\epsilon}$ and $E'_{0,\epsilon}$ the hyperplane parallel to $E_{0,\epsilon}$ which contains this point. The simplex (ball) $\bar{\triangle}_{h_\epsilon}(c_\epsilon)$ is the simplex defined by the hyperplanes $E_{0,\epsilon}$ and $E_{1,\epsilon}, \ldots, E_{n,\epsilon}$. These simplices have the following properties:

- For all small ϵ, $\bar{\triangle}_{h_\epsilon}(c_\epsilon)$ does not contain any point of L_ϵ in its interior. As a consequence, $h_\epsilon = h_{\triangle,L_\epsilon}(c_\epsilon)$.

- When $\epsilon \to 0$, the simplices $\bar{\triangle}_{h_\epsilon}(c_\epsilon)$ converge to $\bar{\triangle}_{h_{\triangle,L}(c)}(c)$ (in Gromov-Hausdorff distance for example).

- The point $p_{\epsilon,0}$ is in the interior of the facet $F_{\epsilon,0}$ of the simplex $\triangle_{h_\epsilon}(c_\epsilon)$. In addition for sufficiently small ϵ, each point $p_{\epsilon,i}$ is in the interior of the facet $F_{\epsilon,i}$ of the simplex $\triangle_{h_\epsilon}(c_\epsilon)$. This is true because $\bar{\triangle}_{h_\epsilon}(c_\epsilon) \to \bar{\triangle}_{h_{\triangle,L}(c)}(c)$, $p_{\epsilon,i} \to p_i$, and each point p_i is in the interior of the facet F_i of $\bar{\triangle}_{h_{\triangle,L}(c)}(c)$.

These properties show that the point c_ϵ is critical for L_ϵ and $\lim_{\epsilon\to 0} c_\epsilon = c$, and hence, we know that $\mathrm{Crit}V(O) \subseteq \lim_{\epsilon\to 0}\mathrm{Crit}V_\epsilon(O)$. Since, the lattice L_G is uniform we deduce that every point in $\lim_{\epsilon\to 0}\mathrm{Crit}V_\epsilon(O)$ belongs to $\mathrm{Crit}V(O)$ which completes the proof. □

The proofs of Theorem 2.1.9 and Theorem 2.1.1 are now complete. We note that this representation of c as a limit of c_ϵ is not in general unique. In fact, *Different non-equivalent classes of critical points, up to linear equivalence, can converge in the limit to the same class.*

Theorem 2.1.21. *In the case where all $b_{ij} > 0$, none of the points ν^π for $\pi \in S_n$ is linearly equivalent to another one, i.e., they define different classes in $\mathbb{R}^{n+1}/L$. In particular, the number of different critical points up to linear equivalence is exactly $n!$.*

However for general graphs this number is usually strictly smaller than $n!$, for examples trees just have one point in $\mathrm{Crit}V(O)$ up to linear equivalence.

2.2 Laplacian Lattices of Connected Regular Digraphs

In this section, we briefly describe how to extend partially the results of the previous section to connected regular digraphs. A digraph D is *regular* if the in-degree and out-degree of each vertex are the same. This allows to define a Laplacian matrix for D, almost similar as in the graphic case: if the vertices of D are enumerated by $\{v_0, \ldots, v_n\}$, the matrix representation of the Laplacian D is of the form Equation 2.1 but we do not have symmetry any more. Namely

$$Q = \begin{bmatrix} \delta_0 & -b_{01} & -b_{02}\ldots & -b_{0n} \\ -b_{10} & \delta_1 & -b_{12}\ldots & -b_{1n} \\ \vdots & \vdots & \ddots & \\ -b_{n0} & b_{n1} & -b_{n2}\ldots & \delta_n \end{bmatrix} \tag{2.6}$$

has the following properties:

(C_1) b_{ij}'s are integers and $b_{ij} \geq 0$ for all $0 \leq i \neq j \leq n$.

(C_2) $\delta_i = \sum_{j=1, j\neq i}^n b_{ij} = \sum_{j=1, j\neq i}^n b_{ji}$ (and is the in-degree (= out-degree) of the vertex v_i).

We consider the lattice generated by the rows of the Laplacian matrix, denote it by L. We obtain the simplicial decomposition of H_0 defined by $\{\triangle^\pi + p \,|\, \pi \in S_n \text{ and } p \in L\}$, similar to the case of unoriented graphs. In the case where all the coordinates b_{ij} are strictly positive, we can similarly prove the following results (the proofs remain unchanged):

- For all $\pi \in S_n$ and $p \in L$, the point $\nu^\pi + p$ is extremal (c.f. Corollary 2.1.15).
- For every two points p and q in L, we have $V(p) \cap V(q) \neq \emptyset$ if and only if p and q are adjacent in the simplicial decomposition of H_0 defined by $\{\triangle^\pi + p | \pi \in S_n \& p \in L\}$. In other words, $V(p) \cap V(q) \neq \emptyset$ if and only if there exists $\pi \in S_{n+1}$ such that q is a vertex of $\Delta^\pi + p$ (c.f. Corollary 2.1.18).

- The set of extremal points of $\Sigma^c(L_G)$ consists of all the points $\nu^\pi + p$ for $\pi \in S_{n+1}$ and $p \in L_G$, i.e., $\mathrm{Ext}^c(L_G) = \{ \nu^\pi + p \mid \pi \in S_{n+1} \text{ and } p \in L_G \}$. As a consequence, we have $\mathrm{Ext}(L_G) = \{ \nu^\pi + p + (1, \ldots, 1) \mid \pi \in S_{n+1} \text{ and } p \in L_G \}$. More precisely, we have $\mathrm{Crit}V_\triangle(O) = \pi_0(\{ \nu^\pi \mid \pi \in S_{n+1} \})$. (c.f. Theorem 2.1.9).

- We have $g_{min} = -\max_{\pi \in S_n} \deg(\nu^\pi) - n$ and $g_{max} = -\min_{\pi \in S_n} \deg(\nu^\pi) - n$.

- Riemann-Roch Inequality. Remark 1.5.3 can be applied: for $K = (\delta_0 - 2, \ldots, \delta_n - 2)$, we have for all D,

$$g_{min} - \deg(D) - 1 \ \leq \ r(K - D) - r(D) \leq g_{max} - \deg(D) - 1.$$

In the general case, where some of the b_{ij}'s could be zero, unfortunately the limiting argument does not behave quite well. Indeed, there are examples of regular digraphs for which a point ν^π is not a critical point for L for some $\pi \in S_n$. However as the proof of Lemma 2.1.19 shows, we always have $\mathrm{Crit}(L) \subseteq \lim_{\epsilon \to 0} \mathrm{Crit}(L_\epsilon)$. So it could happen that we lose (strong) reflection invariance. Although we do not know in general if such lattices have any sort of reflection invariance, it is still possible to prove a Riemann-Roch inequality for these lattices by taking the limit of the Riemann-Roch inequalities for the lattices L_ϵ. One point in doing this limiting argument is to extend the definition of the rank function to all the points of $\mathbb{R}^{n+1}$ (and not only for integral points). This new rank-function will have image in $\{-1\} \cup \mathbb{R}_+$ and is continuous on the points where it is strictly positive.

In the general case we have the following results:

- Every point of degree $-\min_{\pi \in S_n} \deg(\nu^\pi)$ among the points ν^π is extremal (by a similar limiting argument as in the graphic case). So we have $g_{max} = -\min_{\pi \in S_n} \deg(\nu^\pi) - n$. In addition, $g_{min} \geq -\max_{\pi \in S_n} \deg(\nu^\pi) - n$. Let $\bar{g}_{min} = -\max_{\pi \in S_n} \deg(\nu^\pi) - n$.

- (Riemann-Roch Inequality.) Taking the limit of the family of inequalities $g^\epsilon_{min} - \deg(D) - 1 \leq r_\epsilon(K_\epsilon - D) - r_\epsilon(D) \leq g^\epsilon_{max} - \deg(D) - 1$, where ϵ goes to zero, we get

$$\bar{g}_{min} - \deg(D) - 3 \ \leq \ r(K - D) - r(D) \leq g_{max} - \deg(D) + 1.$$

Here r_ϵ is the rank function for the lattice L_ϵ. This is because $\lim_{\epsilon \to 0} g^\epsilon_{min} = \bar{g}_{min}$; $\lim_{\epsilon \to 0} g^\epsilon_{max} = g_{max}$; and $r(D) + 1 \geq \lim_{\epsilon \to 0} r_\epsilon(D) \geq r(D) - 1$ for all $D \in \mathbb{R}^{n+1}$.

2.3 Two Dimensional Sublattices of A_2

In this section, we consider full-rank sublattices of A_2. First, we show that all these sublattices are reflection invariant. It follows that:

Theorem 2.3.1. *Every sublattice L of A_2 of rank two is reflection invariant.*

Indeed something quite strong holds in dimension two: every two dimensional sublattice L of A_2 is a Laplacian lattice of some regular digraph on three vertices.

Lemma 2.3.2. *Every full rank sublattice of A_2 is the Laplacian lattice of a regular digraph on three vertices.*

Let $\{e_0, e_1, e_2\}$ be the standard basis of H_0 where $e_0 = (2, -1, -1)$, $e_1 = (-1, 2, -1)$ and $e_2 = (-1, -1, 2)$. Let the linear functional g_0, g_1 and g_2 be defined by taking the scaler product with e_0, e_1, e_2 respectively. So for example for $u = (u_0, u_1, u_2)$, $g_0(u) = 2u_0 - u_1 - u_2$. Let b_0, b_1 be a basis of L and $b_2 = -b_0 - b_1$. Let Q be the matrix having b_0, b_1 and b_2 as its first, second and third row, respectively. For $i = 0, 1, 2$, define the cone C_i to be the set of vectors v such that $g_i(v) \geq 0$ and $g_j(v) \leq 0$ for $j \neq i$. As a direct consequence of the fact that the vectors e_0, e_1 and e_2 are, up to a positive scaling, orthogonal projections of the standard orthogonal vectors onto H_0 we have:

Lemma 2.3.3. *For a sublattice of A_2, the basis b_0, b_1, b_2 is the basis defined by a regular digraph if and only if the following holds: for each i, b_i is in the cone C_i.*

We now turn to the proof of Lemma 2.3.2:

Proof of Lemma 2.3.2. We should show the existence of lattice points $\{b_0, b_1, b_2\}$ such that:

(i) $\{b_0, b_1\}$ is a basis of L;

(ii) $b_0 + b_1 + b_2 = O$;

(iii) b_i is contained in the cone C_i.

First consider a shortest vector b_0 of the lattice and a shortest vector of the lattice b_1 that is linearly independent of b_0. In the geometry of numbers literature [77], the basis $\{b_0, b_1\}$ is called a Gauss-reduced basis and in fact, $\{b_0, b_1\}$ forms a basis of the lattice L. We may now assume that b_0 is contained in one of the cones C_0, C_1 or C_2, and without loss of generality C_0. Indeed if b_0 does not belong to any of these cones then $-b_0$ will belong to one of these cones, and we may replace b_0 by $-b_0$. So we assume that b_0 belongs to C_0. By the properties of the Gauss reduced basis, we also know that the angle between b_0 and b_1 is in the interval $[\frac{\pi}{3}, \frac{2\pi}{3}]$. Since the maximum angle between any two points in C_i is $\frac{\pi}{3}$, b_1 is contained in a cone different from C_0 and $-C_0$. Now, if b_1 is not contained in C_1 or C_2 then $-b_2$ will be in C_1 or C_2, and we can replace b_1 by $-b_1$. Remark that $\{b_0, -b_1\}$ will remain a basis. Hence, we may assume without loss of generality that $B = \{b_0, b_1\}$ is a basis of the lattice such that b_0 is contained in cone C_0 and b_1 is contained in cone C_1.
This means that $b_0 = (b_{00}, b_{01}, b_{02})$ and $b_1 = (b_{10}, b_{11}, b_{12})$, where $b_{01}, b_{02}, b_{10}, b_{12} \leq 0$ and $b_{00} = -b_{01} - b_{02} > 0$ and $b_{11} = -b_{10} - b_{12} > 0$. First, we observe that $-b_3 = b_0 + b_1$ is contained in $C_0 \cup C_1 \cup -C_2$, and if it is in $-C_2$, then we have our set of lattice points $\{b_0, b_1, b_2\}$. We now define a procedure which, by updating the set of vectors b_0, b_1, provides at the end the set of lattice points $\{b_0, b_1, b_2\}$ with properties $(i), (ii)$ and (iii) above. The procedure is defined as follows:

(a) If $b_0 + b_1 \in -C_2$ then stop.

(b) Otherwise, if $b_0 + b_1 \in C_0$ replace b_0 by $b_0 + b_1$ and iterate.

(c) Otherwise, if $b_0 + b_1 \in C_1$ replace b_1 by $b_0 + b_1$ and iterate.

(d) Output $\{b_0, b_1, b_2\}$, where $b_2 = -b_0 - b_1$.

We will show that the number of iterations is finite. And this shows that the final output has the desired properties. Indeed, at each iteration $\{b_0, b_1\}$ form a basis of L (if $\{b_0, b_1\}$ is a basis of L then $\{b_0 + b_1, b_1\}$ and $\{b_0, b_0 + b_1\}$ will also be a basis of L), and so by the definition of the procedure, the finiteness of the number of steps shows that at the end we should have $b_0 + b_1 \in -C_2$. To show that the procedure terminates after a finite number of iterations, consider a step of the algorithm: if the step (b) in the procedure happens, then $b_0 + b_1$ should be in C_0 and not in $-C_2$. This means that $0 > b_{01} + b_{11}$, which implies that $|g_1(b_0 + b_1)| < |g_1(b_0)|$. Indeed $g_1(b_0 + b_1) = 3b_{01} + 3b_{11} < 0$ and so $|g_1(b_0 + b_1)| = -3b_{01} - 3b_{11} < -3b_{01} = |g_1(b_0)|$. Furthermore, we have, $0 \leq g_0(b_0 + b_1) \leq g_0(b_0)$, since $g_0(b_1) \leq 0$.

Similarly, if the step (c) in the above procedure happens, then $b_0 + b_1$ should be in C_1 and not in $-C_2$. Hence, we should have $|g_0(b_0 + b_1)| < |g_0(b_1)|$ and $0 \leq g_1(b_0 + b_1) \leq g_1(b_1)$. We infer that, starting form b_0 and b_1, at each iteration one of the two inequalities $|g_1(p)| < |g_1(b_0)|$ or $|g_0(p)| < |g_0(b_1)|$ for $p = b_0 + b_1$ should be satisfied. Furthermore, at every iteration we have $|g_0(p)| \leq |g_0(b_0)|$ and $|g_0(p)| \leq |g_0(b_1)|$. Hence, an upper bound on the number of iterations is the number of lattice points p in C_0 with $|g_0(p)| \leq |g_0(b_0)|$ plus the number of lattice points q in C_1 with $|g_1(q)| \leq |g_1(b_1)|$ and this is indeed finite. $\square$

Remark 2.3.4. In higher dimensions, the analogue of Lemma 2.3.2 is unlikely to be true since a simple calculation shows that the minimum angle between cones C_i and C_j is at least $\pi/3$ (here, as in dimension two $e_0, \ldots, e_n$ is the corresponding basis of H_0 where $e_0 = (n, -1, \ldots, -1), \ldots, e_n = (-1, \ldots, -1, n)$, and g_i is the linear form defined by taking the scaler product with e_i). Indeed, let $p = (\sum_{i\neq 0} p_i, -p_1, \ldots, -p_n) \in C_0 - \{O\}$ and $q = (-q_0, \sum_{i\neq 1} q_i, -q_2, \ldots, -q_n) \in C_1 - \{O\}$. We have

$$\begin{aligned}
\frac{p \cdot q}{|p|_{\ell_2}|q|_{\ell_2}} &= \frac{-\sum_{i\neq 0} p_i q_0 - \sum_{i\neq 1} q_i p_1 + p_2 q_2 + \cdots + p_n q_n}{|p|_{\ell_2}|q|_{\ell_2}} \\
&\leq \frac{p_2 q_2 + \cdots + p_n q_n}{|p|_{\ell_2}|q|_{\ell_2}} \\
&\leq \frac{p_2 q_2 + \cdots + p_n q_n}{2\sqrt{p_2^2 + \cdots + p_n^2}\sqrt{q_2^2 + \cdots + q_n^2}} \leq \frac{1}{2}.
\end{aligned}$$

The two inequalities of the last line follow from the set of inequalities

$$|p|_{\ell_2} = \sqrt{(p_1 + \cdots + p_n)^2 + p_1^2 + \cdots + p_n^2} \geq \sqrt{2(p_1^2 + \cdots + p_n^2)} \geq \sqrt{2(p_2^2 + \cdots + p_n^2)}$$

$$|q|_{\ell_2} \geq \sqrt{q_2^2 + \cdots + q_n^2} \quad \text{(similarly as above)},$$

and the Cauchy-Schwartz inequality. Hence, if the lattice L is generated by a regular digraph, then there exists a basis such that the pairwise angles between the elements of the basis is at least $\frac{\pi}{3}$. But, it is known that there exist lattices that are not weakly orthogonal, see [68]. However, note that the notion of a weakly-orthogonal lattice seems to be slightly different from the notion of a digraphical lattice.

We now characterize all the sublattices of A_2 which are strongly reflection invariant.

Theorem 2.3.5. *A sublattice L of A_2 is strongly reflection invariant if and only if there are two different classes of critical points up to linear equivalence or L is defined by a multi-tree on three vertices (i.e., a graph obtained from a tree by replacing each edge by multiple parallel edges).*

Proof. Let $\{b_0, b_1\}$ be the regular digraph basis of L and $b_2 = -b_0 - b_1$. We consider the triangulation $\{\triangle^\pi + p\}$ of H_0 defined by this basis. Let T be the triangle defined by the convex hull of $\{0, b_0, b_0 + b_1\}$ ($= \triangle^\pi$) and let $\bar{T}$ be the opposite of T, the triangle defined by the convex hull of $\{0, b_2, b_1 + b_2\}$ ($= \triangle^{\bar{\pi}}$), and let c_T and $c_{\bar{T}}$ be $\pi_0(\nu^\pi)$ and $\pi_0(\nu^{\bar{\pi}})$. At least one of the points c_T or $c_{\bar{T}}$ is critical. And in addition the set of critical points of $\mathrm{Crit}V(O)$ is a subset of $\{c_T, c_T - b_0, c_T + b_2, c_{\bar{T}}, c_{\bar{T}} + b_0, c_{\bar{T}} - b_2\}$.

($\Rightarrow$) If c_T and $c_{\bar{T}}$ are both critical points and they are different, we have $\mathrm{Crit}V(O) = \{c_T, c_T - b_0, c_T + b_2, c_{\bar{T}}, c_{\bar{T}} + b_0, c_{\bar{T}} - b_2\}$ and we can directly see that $-\mathrm{Crit}V(O) = \mathrm{Crit}V(O) + t$ where $t = c_T + c_{\bar{T}}$. We verify that the only case when c_T and $c_{\bar{T}}$ are equivalent is when $b_0 = (a, 0, -a)$ and $b_1 = (0, b, -b)$ for $a, b > 0$ (in which case $c_T = \pi_0((0, 0, -a - b)$ and $c_{\bar{T}} = \pi_0((-a, -b, 0))$, and so $c_{\bar{T}} - c_T = b_2$ and the lattice is also uniform). In this case, we also have $\mathrm{Crit}V(O) = \{c_T, c_T - b_0, c_T + b_2, c_{\bar{T}}, c_{\bar{T}} + b_0, c_{\bar{T}} - b_2\}$ and so again $-\mathrm{Crit}V(O) = \mathrm{Crit}V(O) + t$ where $t = -c_T - c_{\bar{T}}$.

($\Leftarrow$) If there is just one critical point up to linear equivalence, let us assume without loss of generality that the critical point is c_T. In this case, $\mathrm{Crit}V(O) = \{c_T, c_T - b_0, c + b_2\}$. We now verify that for any bijection ϕ of $\mathrm{Crit}V(O)$ onto itself, $x + \phi(x)$ cannot be the same over all x in $\mathrm{Crit}V(O)$. □

We end this section by providing an example of a sublattice L of A_2 which is not strongly reflection invariant. By the previous theorem, L should contain only one critical point up to linear equivalence and should not be a multi-tree. (In particular, since we only have one class of critical points, L is uniform and reflection invariant. Hence, it satisfies the Riemann-Roch theorem.)

Consider the rank two sublattice L of A_2 defined by the vectors $b_0 = (7, -7, 0)$ and $b_1 = (-3, 11, -8)$, and let $b_2 = -b_0 - b_1 = (-4, -4, 8)$. These vectors form the rows of the 3×3 matrix Q (which is the Laplacian matrix of a regular digraph).

$$Q = \begin{bmatrix} 7 & -7 & 0 \\ -3 & 11 & -8 \\ -4 & -4 & 8 \end{bmatrix} \tag{2.7}$$

Let π and $\bar{\pi}$ be the permutation corresponding to the order $0 <_\pi 1 <_\pi 2$ and its opposite $1 <_{\bar{\pi}} 0 <_{\bar{\pi}} 2$ as in the proof of Theorem 2.3.5. We have $\nu^\pi = \bigoplus\{b_0, b_0 +$

$b_1, O\} = (0, -7, -8)$ and $\nu^{\bar{\pi}} = \bigoplus\{b_1, b_1 + b_0, O\} = (-3, 0, -8)$. We claim that $\nu^{\bar{\pi}}$ is not an extremal point of $\Sigma^c(L)$ and so $\pi_0(\nu^{\bar{\pi}})$ is not critical. This is true because $\nu^{\bar{\pi}} + (7, -7, 0) = (4, -7, -8) \geq \nu^{\pi}$, and so $\nu^{\bar{\pi}}$ cannot be extremal. The following lemma shows that L cannot have a multi-tree basis, more generally that L is not a graphical lattice i.e., it is not generated by the Laplacian of an undirected connected graph.

Lemma 2.3.6. *The lattice L is not a graphical lattice.*

Proof. We know that $|Crit(L)/L| = 1$. By Lemma 2.3.5, we know that if there exists an undirected connected graph G such that $L_G = L$, then it must be a multi-tree. Let us denote such a tree $T(a, b, i)$ where a and b are the number of different multiedges and i is the label of the vertex with degree $a + b$. We now enumerate the different forms of the $L_{T(a,b,i)}$ and verify that they cannot generate L.

$$L_{T_{a,b,1}} = \begin{bmatrix} a+b & -a & -b \\ -a & a & 0 \\ -b & 0 & b \end{bmatrix} \tag{2.8}$$

In this case the last two rows of $L_{T_{a,b,1}}$ are generated by the first two rows of Q (Equation 2.7) by a matrix of the form:

$$M_1 = \begin{bmatrix} -\lambda_1 & 0 \\ 11\lambda_1 & 7\lambda_2 \end{bmatrix} \tag{2.9}$$

where λ_1, λ_2 are non-zero integers. We now verify that $det(M_1) = -7\lambda_1\lambda_2 \neq \pm 1$.

$$L_{T_{a,b,2}} = \begin{bmatrix} a & -a & 0 \\ -a & a+b & -b \\ 0 & -b & b \end{bmatrix} \tag{2.10}$$

In this case, the first and third rows of $L_{T_{a,b,2}}$ are generated by the first two rows of L by a matrix of the form:

$$M_2 = \begin{bmatrix} \lambda_1 & 0 \\ 3\lambda_2 & 7\lambda_2 \end{bmatrix} \tag{2.11}$$

where λ_1, λ_2 are non-zero integers. We now verify that $det(M_2) = 7\lambda_1\lambda_2 \neq \pm 1$.

$$L_{T_{a,b,3}} = \begin{bmatrix} a & 0 & -a \\ 0 & b & -b \\ -a & -b & a+b \end{bmatrix} \tag{2.12}$$

In this case, the first two rows of $L_{T_{a,b,2}}$ are generated by the first two rows of L by a matrix of the form:

$$M_3 = \begin{bmatrix} 11\lambda_1 & 7\lambda_1 \\ 3\lambda_2 & 7\lambda_2 \end{bmatrix} \tag{2.13}$$

where λ_1, λ_2 are non-zero integers. We now verify that $det(M_2) = 56\lambda_1\lambda_2 \neq \pm 1$.

Now we observe that matrix M_i does not depend on a, b and hence the argument for $T_{b,a,i}$ is the same as that of $T_{a,b,i}$. This concludes the proof.

□

Indeed the above example can be turned into a generic class of examples of non-graphical lattices that are uniform and reflection invariant, that we now explain. Consider a lattice defined by generators of the form $b_0 = (\alpha, -\alpha, 0)$ and $b_1 = (-\gamma, \gamma+\eta, -\eta)$:

$$Q = \begin{bmatrix} \alpha & -\alpha & 0 \\ -\gamma & \gamma+\eta & -\eta \\ \gamma-\alpha & -\gamma+\alpha-\eta & \eta \end{bmatrix} \tag{2.14}$$

Here we suppose in addition that $\alpha, \gamma, \eta > 0$ and $\gamma < \alpha \leq \eta + \gamma$ such that the above matrix is the Laplacian of a regular digraph. The two permutations π and $\bar{\pi}$ are defined as above, so for these permutations we have $\nu^{\pi} = (0, -\alpha, -\eta)$ and $\nu^{\bar{\pi}} = (-\gamma, 0, -\eta)$. It is clear that $\deg(\nu^{\pi}) < \deg(\nu^{\bar{\pi}})$. We infer that ν^{π} is extremal. But $\nu^{\bar{\pi}}$ is not extremal since $\nu^{\bar{\pi}} \geq \nu^{\pi} - b_0$.

2.4 Riemann-Roch lattices that are not Graphical.

In this section, we construct an infinite family of sublattices $\{L_n\}_{n=2}^{\infty}$, where L_n is a full rank sublattice of A_n, each L_n satisfies the Riemann-Roch theorem (we say that it has the Riemann-Roch property), and such that none of L_n is graphical. By not being graphical, we mean that there does not exist any basis of L which comes from a connected unoriented multi-graph, i.e. $L_n \neq L_G$ for any connected multi-graph G on $n+1$ vertices.

Indeed, we have already provided in the previous section such an example (and even an infinite number of them) in dimension two: the family of sublattices of A_2 defined by $b_0 = (\alpha, -\alpha, 0)$ and $b_1 = (-\gamma, \gamma+\eta, -\eta)$ (we will prove this shortly below). The construction of L_n for larger values of n is then recursive. Suppose we have already constructed an infinite family of full rank sublattices of A_n which are not graphical and have the Riemann-Roch property, and let L_n be an element of this family. Then we construct a full rank sublattice of A_{n+1} as follows. By taking the natural embedding $A_n \subset A_{n+1}$, $(x_0, \dots, x_n) \to (x_0, \dots, x_n, 0)$, we embed L_n in A_{n+1}. The lattice L_{n+1} is obtained by adding $b_n = (0, 0, \dots, 0, -1, 1)$ to the image of L_n. Remark that if L_n comes from a regular digraph G with vertices $v_0, \dots, v_n$, then L_{n+1} is the lattice of the digraph G' consisting of G and a new vertex v_{n+1} which is connected to v_n by two arcs, one in each direction. We will see that L_{n+1} will not be graphical, and in addition it will have the Riemann-Roch property. Here we provide the details of the construction.

2.4.1 The Lattices L_2

Let L_2 be a sublattice of A_2 defined by $b_0 = (\alpha, -\alpha, 0)$ and $b_1 = (-\gamma, \gamma+\eta, -\eta)$, where $\alpha, \gamma, \eta > 0$ and $\gamma < \alpha \leq \eta + \gamma$.

Proposition 2.4.1. *The sublattice L_2 has Riemann-Roch property and L_2 is not graphical.*

Proof. We saw in the previous section that L_2 has only one class of critical points, up to linear equivalence, is not strongly reflection invariant, and in addition $|\mathrm{Crit}V(O)| = 3$. This shows that L_2 cannot be graphical. However, L_2 is uniform and reflection invariant, and so it has the Riemann-Roch property. $\square$

2.4.2 The Lattices L_n

Let L_n be a full rank sublattice of A_n that we regard as an n-dimensional sublattice of A_{n+1} by taking the embedding $A_n \subset A_{n+1}$ described above. Define L_{n+1} to be the lattice generated by L_n and $b_{n+1} = (0, \ldots, 0, -1, 1)$. We first provide two correspondences: one between the rank function r_n of L_n and the rank function r_{n+1} of L_{n+1}, and the other one, between the extremal points of L_n and the extremal points of L_{n+1}.

Let D be an element of $\mathbb{Z}^{n+2}$. By $D|_n$ we denote the projection of D to $\mathbb{Z}^{n+1}$ obtained by eliminating the last coordinate. So if $D = (D_0, \ldots, D_{n+1})$, then $D|_n = (D_0, \ldots, D_n)$.

Lemma 2.4.2. *Let $D = (D_0, \ldots, D_{n+1})$ be a point in $\mathbb{Z}^{n+2}$ and let $D' = (D - D_{n+1}b_{n+1})|_{n+1}$. We have $r_{n+1}(D) = r_n(D')$.*

Proof. We first prove that $r_n(D') \geq r_{n+1}(D)$. Let $E' \in \mathbb{Z}^{n+1}$ be effective. We should prove that if $\deg(E') \leq r_{n+1}(D)$, then $D' - E' \geq q'$ for at least one $q' \in L_n$. Let $E = (E', 0)$. As $\deg(E) \leq r_{n+1}(D)$, there exists a $q \in L_{n+1}$ such that $D - E \geq q$. By the definition of L_{n+1}, there exists $q' \in L$ and $\alpha \in \mathbb{Z}$ such that $q = (q', 0) + \alpha b_{n+1}$. It follows that $D_{n+1} \geq \alpha$, and so $D' - E' = (D - D_{n+1}b_{n+1} - E)|_n \geq (D - \alpha b_{n+1} - E)|_n \geq q'$. So $D' - E' \geq q'$ and we are done.

We now show that $r_{n+1}(D) \geq r_n(D')$. Let $E = (E_0, \ldots, E_{n+1}) \in \mathbb{Z}^{n+2}$ be effective of degree at most $r_n(D')$. We have to prove the existence of a point $q \in L_{n+1}$ such that $D - E \geq q$. Let $O \leq E' \in \mathbb{Z}^{n+1}$ be defined by $E - E_{n+1}b_{n+1} = (E', 0)$. In other words $E' = (E_0, \ldots, E_{n-1}, E_n + E_{n+1})$. It is clear that $E' \geq O$ and $\deg(E') \leq r_n(D')$. So there exists a point $q' \in L_n$ such that $D' - E' \geq q'$. We infer that $D - E \geq (q', 0) + (D_{n+1} + E_{n+1})b_{n+1}$. So for $q = (q', 0) + (D_{n+1} + E_{n+1})b_{n+1} \in L_{n+1}$, we have $D - E \geq q$, and we are done. $\square$

Lemma 2.4.3. *The extremal points of $\Sigma(L_{n+1})$ are of the form $(v, 0) + q$ where v is an extremal point $\Sigma(L_n)$ and q is a point in L_{n+1}. Similarly, the elements of $\mathrm{Ext}^c(L_{n+1})$ are of the form $(u, -1) + q$ where u is a point of $\mathrm{Ext}^c(L_n)$ and $q \in L_{n+1}$.*

Proof. The proof is similar to the proof of the previous lemma and we only prove one direction, namely $\mathrm{Ext}(L_{n+1}) \subseteq \mathrm{Ext}(L_n) \times \{0\} + L_{n+1}$. The other inclusions $\mathrm{Ext}(L_n) \times \{0\} + L_{n+1} \subseteq \mathrm{Ext}(L_{n+1})$, $\mathrm{Ext}^c(L_{n+1}) \subseteq \mathrm{Ext}^c(L_n) \times \{-1\} + L_{n+1}$, and $\mathrm{Ext}^c(L_n) \times \{-1\} + L_{n+1} \subseteq \mathrm{Ext}^c(L_{n+1})$ involves a similar argument as above.

Let $\bar{v} = (\bar{v}_0, \ldots, \bar{v}_{n+1})$ be an extremal point of L_{n+1}, i.e., $\bar{v} \in \mathrm{Ext}(L_{n+1})$. Let $v \in \mathbb{Z}^{n+1}$ be defined as follows: $(v, 0) = \bar{v} - \bar{v}_{n+1}b_{n+1}$. The claim follows once we have shown that

v is an extremal point of L_n. To prove that v is an extremal point, we need to show that for all $q \in L_n$, $v \not\leq q$ and that v is a local minimum for the degree function. Suppose that this is not the case and let $q \in L_n$ be such that $v \leq q$. We have $\bar{v} \leq (q, 0) + \bar{v}_{n+1}b_{n+1}$ and $(q, 0) + \bar{v}_{n+1}b_{n+1} \in L_{n+1}$, which is a contradiction to the assumption that $\bar{v} \in \mathrm{Ext}(L_{n+1})$. The proof that v is a local minimum follows similarly. □

As a corollary to the above lemmas, we obtain

Corollary 2.4.4. *If L_n has the Riemann-Roch property (resp. is uniform and reflection-invariant), then L_{n+1} also has the Riemann-Roch property (resp. is uniform and reflection-invariant). Furthermore, we have $K_{n+1} = (K_n, 0)$, where K_i is canonical for L_i, $i = n, n+1$.*

We now show that if L_2 is the family of lattices that we described above, then L_n is not graphical. By applying Lemma 2.4.3 and by induction on n, we can show that L_n is not strongly reflection invariant, provided that L_2 is not strongly reflection invariant, and we know that this is the case. Remark that the family of all L_n constructed above is infinite (for each fixed n). Indeed, by using the fact that $\mathrm{Pic}(L_n) = \mathrm{Pic}(L_{n+1})$, and by observing that the set $|\mathrm{Pic}(L_2)|$ contains an infinite number of values, we conclude that $|\mathrm{Pic}(L_n)|$ takes an infinite number of values and so the family of all L_n is infinite.

Chapter 3

The Geometry of the Laplacian Lattice

In this chapter, we obtain combinatorial interpretations for various geometric invariants of the Laplacian lattice, namely the norm of the shortest vector, the covering and packing radius, the Voronoi neighbours and Delaunay triangulation. In the next chapter, we will use this understanding of the Delaunay triangulation to answer some natural questions that arise from the correspondence between the graph and its Laplacian lattice.

3.1 Delaunay Triangulations

Recall that for a point $q \in \mathbb{R}^{n+1}$ and $\lambda \geq 0$, we denote the polytope $\lambda \cdot \mathcal{P} + q$ by $\mathcal{P}(q, \lambda)$. We begin by formalising the notion Delaunay triangulation of a point set under a polyhedral distance function as follows:

Definition 3.1.1. (Delaunay Triangulation under a Polyhedral Distance Function) *A triangulation $\mathcal{T}$ of a discrete point set S in $\mathbb{R}^d$ is a Delaunay triangulation of S under the polyhedral distance function $d_{\mathcal{S}}$ if for every point c in $Crit_{\mathcal{P}}(S)$ there exists a simplex K in $\mathcal{T}$ such that $Q(c, h_{\mathcal{P},S}(c))$ contains the vertices of K in its boundary.*

Remark 3.1.2. The Delaunay triangulation under the simplicial distance function $d_{\triangle}$ is closely related to the notion of Scarf complex associated with a lattice [74]. In fact, in the case of multigraphs whose Laplacian lattice has no zero entries, the Scarf complex coincides with the Delaunay triangulation, but in general the Scarf complex is a subcomplex of the Delaunay triangulation. The Scarf complex of a lattice was first considered in the context of mathematical economics and integer programming, and was later used in commutative algebra to determine free resolutions of the associated lattice ideal. In fact, our determination of the Delaunay triangulation can be reinterpreted as the determination of the minimal free resolution of a generic Laplacian lattice ideal. See Sturmfels and Peeva [70] for more details on the connection between the Scarf complex and minimal free resolution of lattice ideals. The author is indebted to Bernd Sturmfels for pointing out this connection to him. In joint work with Bernd Sturmfels is currently investigating this topic.

We now undertake a detailed study of the Laplacian lattice under the simplicial distance function $d_{\triangle}$. Though we do not always mention it explicitly we always assume that the underlying distance is the simplicial distance function $d_{\triangle}$.

3.2 Laplacian lattice and the Simplicial Distance Function

In this section, we describe the Voronoi diagram, Delaunay triangulation and the local maxima of the simplicial distance function induced by the distance function $d_{\triangle}$ on the Laplacian lattice. We draw heavily from the last chapter and in fact, repeat some of the results. This section is intended to serve two purposes, it contributes to providing a complete description of the geometry of the Laplacian lattice under the simplicial distance function $d_{\triangle}$ and secondly, to collect results are frequently used in the rest of this chapter as well as Chapter 4.

Let $\{b_0, \dots, b_n\}$ be the rows of the Laplacian matrix of G. For each permutation $\sigma \in S_{n+1}$, where S_{n+1} is the symmetric group on $n+1$-elements. we define

$$u_i^{\sigma} = \sum_{j=0}^{i-1} b_{\sigma(i)} \tag{3.1}$$

for i from 0 to n. Note that $u_n^{\pi} = O$ for every permutation $\pi \in S_{n+1}$.

3.2.1 Local Maxima of the Simplicial Distance Function

Given a permutation π on the $n+1$ vertices of G, define the ordering $\pi(v_0) <_{\pi} \pi(v_1) <_{\pi} \dots <_{\pi} \pi(v_n)$ and orient the edges of graph G according to the ordering defined by π i.e. there is an oriented edge from v_i to v_j if $(v_i, v_j) \in E$ and if $v_i <_{\pi} v_j$ in the ordering defined by π. Consider the acyclic orientation induced by a permutation π on the set of vertices of G and define $\nu_{\pi} = (indeg_{\pi}(v_0), \dots, indeg_{\pi}(v_n))$, where $indeg_{\pi}(v)$ is the indegree of the vertex v in the directed graph oriented according to π. Define $\mathrm{Ext}^c(L_G) = \{\nu_{\pi} + q |\ \pi \in S_{n+1},\ q \in L_G\}$.

Theorem 3.2.1. (Chapter 2, Theorem 2.1.9, item (ii)) *The elements of the local maxima of the simplicial distance function $h_{\triangle, L_G}$ are precisely the orthogonal projections of the elements of $Ext^c(L_G)$ onto H_0.*

3.2.2 Voronoi Diagram

We shall first consider Laplacian lattices generated by multigraphs whose Laplacian matrix has no zero entries, in other multigraphs where every pair of vertices are connected by an edge (for example, the complete graph). As we already saw in the previous chapter, these graphs are technically easier to handle than the general case. Recall that two distinct lattice points p and q are Voronoi neighbours if the intersection of their Voronoi cells are non-empty. From Corollary 2.1.18 of Chapter 2, have:

Theorem 3.2.2. *For the Laplacian lattice of a multigraph with no zero entries in its Laplacian matrix, a lattice point q is a Voronoi neighbour of the origin with respect to $d_{\triangle}$ if and only if q is of the form u_i^{σ} for some $\sigma \in S_{n+1}$ and an integer i from 1 to $n-1$.*

The case of general multigraphs is slightly more involved. Since Theorem 3.2.3 is not used in the rest of the thesis, the reader may skip the theorem for the first reading.

Theorem 3.2.3. *The Voronoi neighbours of the origin with respect to $d_{\triangle}$ are precisely the set of non-zero lattice points that are contained in $\bar{\triangle}(c_{\pi}, h_{\triangle,L_G}(c_{\pi}))$ where $c_{\pi} = \pi_0(\nu_{\pi})$.*

Proof. First, we know that the Voronoi cell of every non-zero lattice point q in $\bar{\triangle}(c_{\pi}, h_{\triangle,L_G}(c_{\pi}))$ shares c_{π} with the Voronoi cell of the origin and hence q is a Voronoi neighbour of the origin. Conversely, consider a lattice point q that is a Voronoi neighbour of the origin O we know that the simplices centered at q and O share a point m say in the boundary of the arrangement of simplices $\triangle$ centered at lattice points and the radius of simplices is r for some real number $r > 0$. By the duality theorem (see Appendix, Theorem 7.0.12) we know that there is a point c in $\mathrm{Crit}(L_G)$ such that $\bar{\triangle}(c, \mathrm{Cov}_{\triangle}(L_G)-r)$ contains m. Applying triangle inequality, we deduce that $d_{\bar{\triangle}}(c, O) \leq \mathrm{Cov}_{\triangle}(L_G)$ but we also know that $d_{\bar{\triangle}}(c, q') \geq \mathrm{Cov}_{\triangle}(L_G)$ for all $q' \in L_G$. Hence, $d_{\bar{\triangle}}(c, O) = d_{\triangle}(O, c) = \mathrm{Cov}_{\triangle}(L_G)$. By item ii of Theorem 2.1.9 we know that $c = c_{\pi}$ for some permutation $\pi \in S_{n+1}$. Similarly, we also know that $d_{\bar{\triangle}}(c, q) = \mathrm{Cov}_{\triangle}(L_G)$ and hence, q is contained in $\bar{\triangle}(c_{\pi}, h_{\triangle,L_G}(c_{\pi}))$. $\square$

For the case of multigraphs with no zero entries in its Laplacian matrix, Theorem 3.2.2 gives a useful characterization of the Voronoi neighbours of the origin, but for the case of general connected graphs, the characterization obtained in Theorem 3.2.3 is not explicit enough for our purposes. In the subsequent sections, we use the following perturbation trick that we also used in Chapter 2 to handle the Laplacian lattice of a general connected graph: we perturb the Laplacian lattice and scale it to the "nice" case, i.e. to the case of lattices generated by multigraphs with no zero entries in its Laplacian matrix, and study the limit as the perturbation tends to zero. More precisely, we consider lattices generated by the following perturbed basis: we add a rational number $\epsilon > 0$ to every non-diagonal element b_{ij} of the Laplacian matrix and then set the diagonal elements so that the row sum and column sum is zero. We call such a perturbation a **standard** perturbation and denote the vector obtained by perturbing b_i by b_i^{ϵ}. The following lemma characterizes the Voronoi neighbours of the perturbed Laplacian lattice:

Lemma 3.2.4. *Let L_G be the Laplacian lattice of the graph and L_G^{ϵ} be the lattice obtained by perturbing L_G according to the standard perturbation. The Voronoi neighbours of L^{ϵ} under the distance function $d_{\triangle}$ are of the form $u_S^{\epsilon} = b_{i_1}^{\epsilon} + b_{i_2}^{\epsilon} + \cdots + b_{i_k}^{\epsilon}$ where i_js are distinct.*

Proof. Since the perturbation ϵ is rational, we can scale the lattice by a factor λ, say to obtain the Laplacian lattice of a multigraph with no zero entries in its Laplacian matrix. We can apply Theorem 3.2.2 to the scaled lattice λL^{ϵ} to deduce that the Voronoi neighbours of the lattice λL^{ϵ} are of the form λu_S^{ϵ} for some subset S. Finally, observe that the Voronoi neighbours of a lattice are preserved under scaling to complete the proof. □

Remark 3.2.5. Note that there are different ways of perturbing the basis to obtain a "nice" lattice. For example, we can add an $\epsilon > 0$ only to the non-diagonal elements that are zero and then set the diagonal elements such that the row sum and column sum is zero.

3.2.3 Delaunay Triangulation

From the study of Laplacian lattices in Chapter 2 we have:

Theorem 3.2.6. *Let S^{σ} be the convex hull of $u_0^{\sigma}, \dots, u_n^{\sigma}$. The set $\{S^{\sigma} + p\}_{\sigma \in S_{n+1}, p \in L_G}$ is a Delaunay triangulation of L_G and is a unique Delaunay triangulation if G is a multigraph with no zero entries in its Laplacian matrix.*

We shall see in Section 4.2 that the number of different graphs that have L_G as their Laplacian lattice upper bounds the number of "different" Delaunay triangulations of L_G.

3.3 Packing and Covering Radius of the Laplacian Lattice

We will show that the packing radius of the graph under the distance function $\triangle$ (and $\bar{\triangle}$) is essentially (up to a factor depending on the number of vertices) the minimum cut of a graph.

Definition 3.3.1. **(ℓ_1-Minimum Cut)** *For a non-trivial cut S of $V(G)$, define the weight of the cut $\mu_1(S) = \sum_{v \in S} deg_{S,\bar{S}}(v)$ where $deg_{S,\bar{S}}(v)$ is the degree of the vertex v across the cut S. Now define the ℓ_1-minimum cut $MC_1(G)$ as the minimum of $\mu_1(S)$ over all non-trivial cuts S i.e. S is neither empty nor equal to $V(G)$.*

Remark that ℓ_1-minimum cut of a graph is the same as the minimum cut of a graph. We call it the ℓ_1-minimum cut to distinguish from a variant the "ℓ_{∞}-minimum cut" that we will encounter in the next section. Recall that for points $p, q \in \mathbb{R}^{n+1}$, the tropical sum $p \oplus q$ is defined as $(min(p_1, q_1), \dots, min(p_{n+1}, q_{n+1}))$.

Lemma 3.3.2. *For any point p in H_0, $\triangle$-midpoint m of p and the origin O is the projection of the max-sum of the two points onto H_0 and $d_{\triangle}(p, m) = d_{\triangle}(O, m) = || - p \oplus O||_1/(n+1)$.*

Proof. For a point $r \in \mathbb{R}^{n+1}$ let H_r^+ be the domination cone defined as: $H_r^+ = \{r' | \ r' \in \mathbb{R}^{n+1}, \ r' - r \geq O\}$. Consider H_p^+ and H_O^+ and observe that the point closest to H_0 in their intersection is $-(-p \oplus O)$. Project the system onto H_0 along the normal $(1, \ldots, 1)$. A simple computation shows that the projection of the cones onto H_0 are simplices and these simplices are dilated and translated copies of $\triangle$ and the only point of intersection of the simplices is the projection of the max-sum. Hence, we obtain $d_\triangle(O, m) = d_\triangle(p, m) = || - p \oplus O||_1/(n+1)$. □

Lemma 3.3.3. *For a subset S of the rows of the Laplacian matrix of the graph, let $u_S = \sum_{b_i \in S} b_i$. The ℓ_1-norm of the max-sum of u_S and O is the size of the cut defined by S.*

Proof. The max sum of u_S and O is given by:

$$(-u_S \oplus O)_i = \begin{cases} -(u_S)_i, \ \text{if } (u_S)_i > 0, \\ 0, \ \text{otherwise.} \end{cases} \tag{3.2}$$

Hence, we consider the positive coordinates of u_S. Now note that since the sum of coordinates of u_S is zero, the absolute sum of the positive coordinate is equal to absolute sum of negative coordinates. Furthermore, the negative coordinates are characterized by vertices that do not belong to the cut and the sum of the absolute values of the negative-valued coordinates of u_S is the size of the cut S. □

Theorem 3.3.4. *The packing radius of L_G under the simplicial distance function $d_\triangle$ is equal to $\frac{MC_1(G)}{n+1}$ where $MC_1(G)$ is the size of the ℓ_1-minimum cut.*

Proof. First, observe that the lattice point that defines the packing radius under the distance function $d_\triangle$ is a Voronoi neighbour of the origin under $d_\triangle$. Hence, we only restrict to the Voronoi neighbours. Now for the case of general connected graphs, the characterization of the Voronoi neighbours of the origin obtained in Theorem 3.2.3 is not explicit enough. Hence, we perform the standard perturbation of the Laplacian matrix (see Subsection 3.2.2) and consider lattices generated by these perturbed matrices. Using Lemma 3.2.4, we know that the packing radius of the perturbed lattice L_G^ϵ is defined by a point of the form $u_S = \sum_{i \in S} b_i^\epsilon$ for some non-trivial subset S of $V(G)$ and b_i^ϵ being the perturbed vector of b_i. Using the fact that the packing radius is preserved under perturbation (see Appendix Lemma 7.0.13), we deduce that the packing radius is defined by a point of the form u_S for a non-trivial subset S of $V(G)$. By Lemma 3.3.2 and the definition of packing radius we have $\text{Pac}_\triangle(L_G) = \min_S || - u_S \oplus O||_1/n + 1$ where the minimum is taken over all the non-trivial cuts S. We now use Lemma 3.3.3 to deduce that $\min_S || - u_S \oplus O||_1$ is equal to the size of the minimum cut of G. □

Since we know that the Laplacian lattice is uniform and every extremal point has degree equal to g, the genus of the graph we have the following theorem.

Theorem 3.3.5. *The covering radius of the Laplacian lattice is equal to $\frac{g+n}{n+1}$.*

3.4 The Shortest Vector of the Laplacian Lattice

We will provide a combinatorial interpretation of the norm of the shortest vector of the Laplacian lattice under the simplicial distance function $d_\triangle$. We will see that the length of the shortest vector of the Laplacian lattice is in fact a certain variant of the minimum cut in the graph. A precise definition follows:

Definition 3.4.1. **(ℓ_∞-Minimum Cut)** *For a non-trivial cut S of $V(G)$, define the weight of the cut as $\mu_\infty(S) = \max\{deg_{S,\bar{S}}(v) |\ v \in \bar{S}\}$ where $deg_{S,\bar{S}}(v)$ is the degree of the vertex v across the cut S. Now define the ℓ_∞-minimum cut $MC_\infty(G)$ as the minimum of $\mu_\infty(S)$ over all non-trivial cuts S i.e. S is neither empty nor equal to $V(G)$.*

Note that for a simple connected graph G, we have $\mathrm{MC}_\infty(G) = 1$.

Theorem 3.4.2. *The length of the shortest vector $\nu_\triangle(L_G)$ of the Laplacian lattice under the simplicial distance function $d_\triangle$ is equal to* $\mathrm{MC}_\infty(G)$.

Proof. First let us consider the case where G is a multigraph with no zero entries in its Laplacian matrix. By Lemma 1.4.8, we know that every shortest vector in the distance function $d_\triangle$ must be a Voronoi neighbour of the origin under $d_\triangle$. We know that the Voronoi neighbours of the origin under $d_\triangle$ are of the form: $u_S = \sum_{i \in S} b_i$ for some non-trivial subset S of $V(G)$. For the case of general connected graphs, the characterization of the Voronoi neighbours of the origin obtained in Theorem 3.2.3 is not explicit enough. Hence, we perform the standard perturbation (see Subsection 3.2.2) of the Laplacian matrix and consider lattices generated by these perturbed matrices. By Lemma 3.2.4, we know that the shortest vector of the perturbed lattice L^ϵ_G is defined by a point of the form $u_S = \sum_{i \in S} b^\epsilon_i$ for some non-trivial subset S of $V(G)$ and b^ϵ_i being the perturbed vector of b_i. Using the fact that the quantity $\nu_\triangle(.)$ is preserved as the perturbation tends to zero (See Appendix Lemma 7.0.13), we deduce that the shortest vector for the general lattice must be of the form: $u_S = \sum_{i \in S} b_i$ for some non-trivial subset S of $V(G)$.

Consider a subset S such that u_S is a shortest vector. First, recall that $d_\triangle(O, u_S) = |\min_j\ u_{Sj}|$. Now, since $u_S \neq O$ only the negative coordinates of u_S define $d_\triangle(O, u_S)$ and the negative coordinates are indices j such that $v_j \notin S$. Hence $d_\triangle(O, u_S) = \max\{deg_{S,\bar{S}}(v)\ |v \in \bar{S}\}$. We have:

$$\nu_\triangle(L_G) = \min_S\{d_\triangle(O, u_S)\} = \min_S \max\{deg_{S,\bar{S}}(v)\ |v \in \bar{S}\} = \mathrm{MC}_\infty(G). \tag{3.3}$$

□

Corollary 3.4.3. *For a simple connected graph G, we have $\nu_\triangle(L_G) = 1$.*

Chapter 4

Applications of the Geometric Study

In this chapter, we exploit our understanding of the Laplacian lattice to answer some natural questions that arise from the correspondence between the graph and its Laplacian lattice.

4.1 The Delaunay Polytope of the Laplacian Lattice

A natural question that arises with the correspondence between the Laplacian lattice and a graph is whether the Laplacian lattice characterizes the underlying graph completely up to isomorphism? The first observation towards answering this question are the following lemmas:

Lemma 4.1.1. *[13] The covolume of the Laplacian lattice of G with respect to A_n is equal to the number of spanning trees of G.*

Lemma 4.1.2. *The Laplacian lattice of any tree on $n+1$ vertices is the root lattice A_n.*

Proof. We know from Lemma 4.1.1 that the Laplacian lattice is a sublattice of the root lattice A_n and the covolume of a Laplacian lattice with respect to A_n is equal to the number of spanning trees of the graph. This implies that in the case of trees the covolume of the Laplacian lattice is equal to one. Hence, the Laplacian lattice of a tree is the root lattice A_n itself. □

Lemma 4.1.2 shows that the Laplacian lattice itself does not characterize a graph completely up to isomorphism. However, we will now see that the Delaunay triangulations of the Laplacian lattice L with respect to the simplicial distance function provide more refined information about the graphs that have L as their Laplacian lattice. More precisely, each graph provides a Delaunay triangulation of its Laplacian lattice and the Delaunay polytope of the origin characterizes the graph completely up to isomorphism. In the course of showing this result, we also study the structure of the Delaunay polytope in particular determine its vertices, facets and edges. Recall that as we noted in Remark 3.1.2, in the case of graphs with no zero entries in the Laplacian matrix our study of the

Delaunay triangulation is in fact a study of the Scarf complex of the Laplacian lattice and in general, the Delaunay triangulation contains the Scarf complex of L_G.

4.1.1 The Structure of the Delaunay Polytope

We mainly rely on the following simple facts to study the structure of the Delaunay polytope:

1. The Separation theorem for closed convex sets [75]: two non-empty disjoint compact convex sets can be separated by a hyperplane. Equivalently, for two non-empty disjoint compact convex sets C_1, C_2, there is a linear functional f such that $f(x) < 0$ for all $x \in C_1$ and $f(y) > 0$ for all $y \in C_2$.

2. The Laplacian matrix of an undirected graph is symmetric.

3. If G is a connected graph on $n+1$-vertices then the Laplacian matrix $Q(G)$ has rank n and hence, its Laplacian lattice has dimension n.

We know from Theorem 3.2.6 that each graph provides a Delaunay triangulation of its Laplacian lattice in the following manner: Recall that $\{b_0, \dots, b_n\}$ are the rows of the Laplacian matrix of G. For each permutation $\sigma \in S_{n+1}$, we define $u_i^\sigma = \sum_{j=0}^{i} b_{\sigma(j)}$ for i from 0 to n. Note that $u_n^\pi = (0, \dots, 0)$ for every permutation $\pi \in S_{n+1}$. We define the simplex $\triangle_\sigma$ as the convex hull of $u_0^\sigma, \dots, u_n^\sigma$. The Delaunay polytope of the origin i.e. the set of Delaunay simplices with the origin as a vertex, is given by $\{\triangle_\sigma\}_{\sigma \in S_{n+1}}$. See Figure 4.1 for the Delaunay polytope of some small graphs. We denote this polytope by $H_{Del_G}(O)$.

We first describe the vertex set of $H_{Del_G(O)}$. A study of low-dimensional examples leads us to the claim that every vertex of $H_{Del_G}(O)$ is of the form u_k^σ for some integer k from 0 to $n-1$ and a permutation $\sigma \in S_{n+1}$. In order to show this claim, we proceed as follows: we consider the convex hull $H'(G)$ of points u_k^σ for k from $0, \dots, n$ and $\sigma \in S_{n+1}$. We then show that every point of the form u_k^σ for k from $0, \dots, n-1$ is a vertex of $H'(G)$ and that $H'(G) = H_{Del_G}(O)$.

Lemma 4.1.3. *Every point of the form u_k^σ is a vertex of $H'(G)$ where k varies from 0 to $n-1$ and $\sigma \in S_{n+1}$.*

Proof. We show that for every point of the form $u_k^\sigma = \sum_{i=0}^{k} b_\sigma(i)$ for k from 0 to $n-1$ there exists a point w such that the linear functional $f(x) = w \cdot x^t$ has the property:

$$f(b_{\sigma(i)}) = \begin{cases} > 0 \text{ if } 0 \le i \le k, \\ < 0 \text{ otherwise.} \end{cases}$$

The details of the construction of the functional f are as follows. Consider the set $C_{\sigma,k}$ of points $(p_0, \dots, p_n) \in H_0$ such that

$$p_{\sigma(i)} = \begin{cases} > 0 \text{ if } 0 \le i \le k, \\ < 0 \text{ otherwise.} \end{cases}$$

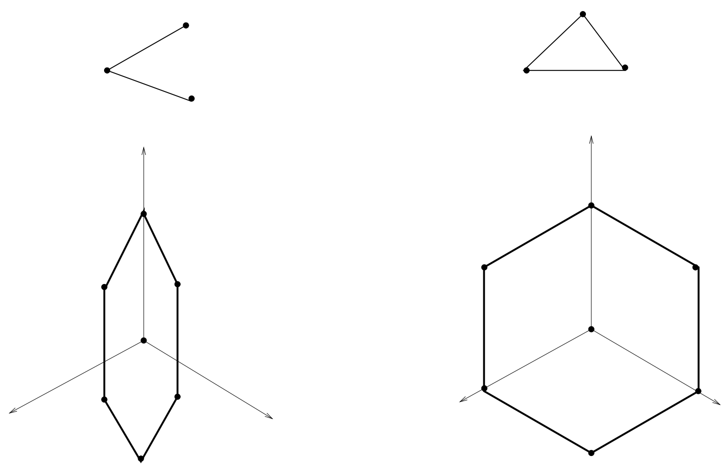

Figure 4.1: The Delaunay polytope of a path on two edges (left) and a triangle.

We make the following observations:

1. $C_{\sigma,k}$ is a cone.
2. $C_{\sigma,k}$ is not empty for k from 0 to $n-1$.

Take a point q in $C_{\sigma,k}$. Since G is connected, L_G is a lattice of dimension n and a basis of L_G spans H_0. Consider the basis $\{b_0, \ldots, b_{n-1}\}$ (the first n rows of $Q(G)$). Let $q = \sum_{i=0}^{n-1} v_i \cdot b_i$. Set $w_i = v_i$ for i from 0 to $n-1$ and $w_n = 0$. Using the symmetry of the Laplacian matrix, we can easily verify that the functional f has the desired properties. By the properties of f, it follows that u_k^σ is the unique maximum of $f(x)$ among vertices of the form u_j^π for an arbitrary permutation $\pi \in S_{n+1}$. This implies that u_k^σ is also the unique maximum of $f(x)$ over the polytope $H'(G)$. Using standard arguments in linear optimization [76], this implies that u_k^σ is a vertex of $H'(G)$. $\square$

Next, we characterize points in the simplex $\triangle_\sigma$.

Lemma 4.1.4. *A point p is contained in $\triangle_\sigma$ if and only if it can be written as $p = \sum_{i=0}^{n} \lambda_i b_{\sigma(i)}$ where $1 = \lambda_0 \geq \lambda_1 \geq \cdots \geq \lambda_n \geq 0$.*

Proof. If p is contained in $\triangle_\sigma$, then we can write: $p = \sum_{i=0}^{n} \mu_i u_i^\sigma$ where $\mu_i \geq 0$ and $\sum_{i=0}^{n} \mu_i = 1$. We can plug in $u_k^\sigma = \sum_{i=0}^{k} b_{\sigma(i)}$ to the equation $p = \sum_{i=0}^{n} \mu_i u_i^\sigma$ to obtain: $p = \sum_{i=0}^{n} \lambda_i b_{\sigma(i)}$ where $\lambda_k = \sum_{i=k}^{n} \mu_i$. Observe that $1 = \lambda_0 \geq \lambda_1 \geq \cdots \geq \lambda_n \geq 0$.

Conversely, suppost that a point can be written as $p = \sum_{i=0}^{n} \lambda_i b_{\sigma(i)}$ where $1 = \lambda_0 \geq \lambda_1 \geq \cdots \geq \lambda_n \geq 0$. We set $\mu_i = \lambda_i - \lambda_{i+1}$ for $0 \leq i \leq n-1$ and $\mu_n = \lambda_n$ and we have $p = \sum_{i=0}^{n} \mu_i u_i^\sigma$. We finally verify the following properties: i. $\mu_i \geq 0$, since $\lambda_i \geq \lambda_{i+1}$ for i from 0 to $n-1$, ii. $\lambda_n \geq 0$ and iii. $\sum_{i=0}^{n} \mu_i = \lambda_0 = 1$. This shows that p is contained in $\triangle_\sigma$. $\square$

Corollary 4.1.5. *The set $H_{Del_G}(O)$ is convex.*

Proof. Consider points p_1 and p_2 in simplices $\triangle_{\sigma_1}$ and $\triangle_{\sigma_2}$ for some $\sigma_1, \sigma_2 \in S_{n+1}$. By Lemma 4.1.4, we can write: $p_1 = \sum_{i=0}^{n} \lambda_i^1 b_{\sigma_1(i)}$ and $p_2 = \sum_{i=0}^{n} \lambda_i^2 b_{\sigma_2(i)}$ where $1 = \lambda_0^i \geq \lambda_1^i \geq \cdots \geq \lambda_n^i \geq 0$ for $i \in \{1, 2\}$. Each point p in the line segment joining p_1 and p_2 can be written as $\mu p_1 + (1-\mu)p_2$ for some $0 \leq \mu \leq 1$. We can write: $p = \sum_{i=0}^{n} \lambda_i b_i$ and there exists a permutation σ such that $1 \geq \lambda_{\sigma(0)} \geq \lambda_{\sigma(2)} \geq \cdots \geq \lambda_{\sigma(n)} \geq 0$. We add $O = (1 - \lambda_{\sigma(0)}) \sum_{i=0}^{n} b_i$ to the right hand side of the equation to obtain a point of the form stated in Lemma 4.1.4. This concludes the proof. □

Lemma 4.1.6. *For any undirected connected graph G, we have $H'(G) = H_{Del_G(O)}$.*

Proof. Consider a point p in $H_{Del_G}(O)$. By definition, p belongs to some simplex of the form $\triangle_\sigma$. Observe that all the vertices of $\triangle_\sigma$ except possibly the origin are contained in $H'(G)$. Now, let $p = \sum_{i=0}^{n-1} \lambda_i u_i^\sigma + \lambda_n \cdot O$ for some $\lambda_i \geq 0$ and $\sum_{i=0}^{n} \lambda_i = 1$. We have $p = \sum_{i=0}^{n-1} \lambda_i u_i^\sigma + \lambda_n \cdot O = \sum_{i=0}^{n-1} \lambda_i u_i^\sigma + \lambda_n(b_0 + \sum_{j=1}^{n} b_j)/2$. This shows that p can be written as a convex combination of points in $H'(G)$ and hence p is contained in $H'(G)$. We show the converse by contradiction. Assume that there exists a point p in $H'(G) \setminus H_{Del_G(O)}$. By Corollary 4.1.5, $H_{Del_G(O)}$ is a convex polytope and hence, a closed subset of H_0 equipped with the Euclidean topology. By the separation theorem for closed convex sets [75], there exists a linear functional $f(x)$ such that $f(p) < 0$ and $f(y) > 0$ for all $y \in H_{Del_G(O)}$. Since u_k^σ is contained in $H_{Del_G}(O)$, we know that $f(u_k^\sigma) > 0$ for all $0 \leq k \leq n$ and $\sigma \in S_{n+1}$. But since, p is contained in $H'(G)$ it can be written as a convex combination of the vertices of $H'(G)$ and the vertices of $H'(G)$ are points of the form u_k^σ for k from 0 to $n-1$ and hence $f(p) > 0$. We obtain a contradiction. □

Corollary 4.1.7. *The set $H_{Del_G}(O)$ is a convex polytope and has $2^{n+1} - 2$ vertices.*

We now describe the facet structure of $H_{Del_G}(O)$. We know that every vertex of $H_{Del_G}(O)$ is of the form:

$$v = \sum_{j=0}^{k} b_{i_j}, \text{for } k = 1 \ldots n \text{ and } b_{i_j}\text{s are all distinct.} \tag{4.1}$$

We define the set V_i for $i = 0 \ldots n$ as the subset of vertices that contain b_i in their representation of the form stated in (4.1). Define the set $F_{i,j} = V_i \setminus V_j$.

Lemma 4.1.8. *For each integer $0 \leq i, j \leq n$ and $i \neq j$, the affine hull of the elements of $F_{i,j}$ is an $n-1$-dimensional affine space.*

Proof. For the sake of convenience, we consider $F_{0,n}$. The set $S = \{b_0, b_0 + b_1, \ldots, b_0 + b_1 + b_2 + \cdots + b_{n-1}\}$ is contained in $F_{0,n}$. Since G is connected, the elements of S are linearly independent and hence their affine hull is an $n-1$-dimensional affine space.

We show that every element of $F_{0,n}$ is contained in the affine hull of S. The affine hull of S consists of elements of the form $\sum_{i=0}^{n} \alpha_i u_i$ where $u_i = \sum_{k=0}^{i} b_k$ and $\sum_{i=0}^{n} \alpha_i = 1$. This information can be written as $\sum_{i=0}^{n} \lambda_i b_i$ where $\lambda_i = \sum_{j=i}^{n} \alpha_j$. This means that $\lambda = T\alpha$ where $\lambda = (\lambda_0, \ldots, \lambda_n)$, $\alpha = (\alpha_0, \ldots, \alpha_n)$ and T is the lower triangular matrix with unit entries. By definition, the elements of $F_{0,n}$ are of the form: $b_0 + b_{i_1} + \cdots + b_{i_k}$

where $i_j \notin \{0, n\}$. Consider an element $u = b_0 + b_{j_1} + \cdots + b_{j_k}$ in $F_{0,n+1}$ and let $S' = \{0, j_1, \ldots, j_k\}$. Set $\lambda' = i_{S'}$, the indicator vector of the set S' and since T is invertible there exists a unique vector $\alpha' = T^{-1}\lambda'$ and we verify that $\sum_{i=0}^{n} \alpha'_i$ is the coefficient of b_0 in u and is hence equal to 1. This shows that u is contained in the affine hull of S. The same argument can be done for any $F_{i,j}$ by replacing 0 by i and n by j. □

Lemma 4.1.9. *For each integer $0 \leq i, j \leq n$ and $i \neq j$, the convex hull of the elements of $F_{i,j}$ is a facet of $H_{Del_G}(O)$. Furthermore, each element of $F_{i,j}$ is in convex position.*

Proof. For the sake of convenience, we consider $F_{0,n}$. By Lemma 4.1.8, the affine hull of the elements of $F_{0,n}$ spans a $n-1$-dimensional space K and we know that $H_{Del_G}(O)$ is a n-dimensional polytope. In order to show that $F_{0,n}$ is a facet of $H_{Del_G}(O)$ we need to show that the affine hull of the elements of $F_{0,n}$ supports $H_{Del_G}(O)$. This can be seen as follows: since G is connected we know that $\{b_0, \ldots, b_{n-1}\}$ is a basis of H_0. Hence, every point in H_0 can be uniquely written as $\sum_{i=0}^{n-1} \alpha_i b_i$ for some $\alpha_i \in \mathbb{R}$ and the hyperplane K is given by $\alpha_0 = 1$. Consider a vertex v of $H_{Del_G}(O)$ not contained in $F_{0,n}$. Let $v = \sum_{j=1}^{k} b_{i_j}$. Since v is not contained in $F_{0,n}$, either b_0 is not contained in the sum representing v or b_n is contained in the sum. If b_0 is not contained in the sum and b_n is contained in the sum then $\alpha_0 = -1$. If b_0 and b_n are contained in the sum, then $\alpha_1 = 0$ and also, if b_0 and b_n are not contained in the sum then $\alpha_0 = 0$. This shows that all the vertices of $H_{Del_G}(O)$ that are not in $F_{0,n-1}$ are strictly contained in K^-, the halfspace $\alpha_0 \leq 1$. This suffices to conclude that $K \cap H_{Del_G}(O) = F_{0,n}$. By Lemma 4.1.3, all the elements of $F_{0,n}$ are in convex position and this concludes the proof. □

Lemma 4.1.10. *The facets of $H_{Del_G}(O)$ are exactly of the form $F_{i,j}$ for $0 \leq i, j \leq n$ and $i \neq j$.*

Proof. By construction, the facets of $H_{Del_G}(O)$ must be contained in the affine hull of the facets of $\triangle_\sigma$. By Lemma 4.1.9, we know that the affine hull of any facet of $\triangle_\sigma$ not containing the origin contains a facet of $H_{Del_G}(O)$. It suffices to show that the affine hull of a facet of $\triangle_\sigma$ containing the origin does not contain a facet of $H_{Del_G}(O)$. Consider a facet F of $\triangle_\sigma$. We may assume, without loss of generality, that F contains all vertices of $\triangle_\sigma$ apart from u_k^σ for some $0 \leq k \leq n-1$. Assume that F is a facet of $H_{Del_G}(O)$. This means that there is an affine function f such that

P1. $f(u_i^\sigma) = c$ for all $i \neq k$.

P2. $f(v) > c$ for all vertices of $H_{Del_G}(O)$ that are not contained in the affine hull of F.

Note that by the property P1, we have $f(O) = c$. Consider the linear function $g(x) = f(x) - f(O)$. We have $g(u_i^\sigma) = 0$ for all $i \neq k$. Now, suppose that $g(u_k^\sigma) = 0$, then $f(u_k^\sigma) = f(u_j^\sigma)$ for all $j \neq k$, but since G is connected, u_k^σ is not contained in the affine hull of F and this contradicts the property P2. If $g(u_{\sigma(k)}) < 0$, then $f(u_k^\sigma) < c$ and hence again contradicts property P2. On the other hand, if $g(u_{\sigma(k)}) > 0$, then $g(-u_{\sigma(k)}) < 0$ and $f(-u_k^\sigma) < c$ and $-u_k^\sigma$ is indeed a vertex of $H_{Del_G}(O)$. This again contradicts the property P2. This concludes the proof. □

Corollary 4.1.11. *The number of facets of $H_{Del_G}(O)$ is $n \cdot (n+1)$.*

We now characterize the edges of $H_{Del_G}(O)$.

Lemma 4.1.12. *The edges of $H_{Del_G}(O)$ are of the form $e_k^\sigma = (u_k^\sigma, u_{k+1}^\sigma)$ for k from 0 to $n-1$ and $\sigma \in S_{n+1}$.*

Proof. The proof is along the same lines as the proof of Lemma 4.1.3. We show that for each e_k^σ, there exists a linear function $f(x) = w \cdot x^t$ such that $f(u_k^\sigma) = f(u_{k+1}^\sigma) = c$ and $f(v) < c$ for all other vertices v of $H_{Del_G}(O)$.

Consider the set $C_{e_k^\sigma}$ of points $(p_1, \ldots, p_{n+1}) \in H_0$ such that

$$p_{\sigma(i)} = \begin{cases} > 0 \text{ if } 0 \leq i \leq k, \\ = 0 \text{ if } i = k+1, \\ < 0, \text{otherwise.} \end{cases}$$

We make the following observations:

1. $C_{e_k^\sigma}$ is a cone.
2. $C_{e_k^\sigma}$ for $0 \leq k \leq n-1$ is not empty.

Take a point q in $C_{\sigma,k}$. Since G is connected, L_G is a lattice of dimension n and a basis of L_G spans H_0. Consider the basis $\{b_0, \ldots, b_{n-1}\}$ (the first n rows of $Q(G)$) and write $q = \sum_{i=0}^{n-1} v_i \cdot b_i$. Set $w_i = v_i$ for i from 0 to $n-1$ and $w_n = 0$.

The functional $f(x) = w \cdot x^t$ attains its maximum precisely on the edge e_k^σ. Using standard arguments in linear optimization [76], this implies that e_k^σ is an edge of $H'(G)$.

It remains to show that there are no other edges of $H_{Del_G}(O)$. By construction of $H_{Del_G}(O)$, an edge of $H_{Del_G}(O)$ is contained on the affine hull of an edge of $\triangle^\sigma$ for some $\sigma \in S_{n+1}$. An edge of $\triangle^\sigma$ is of the form (u_i^σ, u_j^σ) for some $0 \leq i, j \leq n$ and $i \neq j$. We know that $(u_k^\sigma, u_{k+1}^\sigma)$ is an edge for $0 \leq k \leq n-1$ and $\sigma \in S_{n+1}$. Consider an edge $e_{i,j}$ of $\triangle^\sigma$ of the form (u_i^σ, u_j^σ) where $j \neq i+1$ and $i < j$. Assume that the affine hull of $e_{i,j}$ contains an edge of $H_{Del_G}(O)$. This implies that there exists an affine function f such that

P1. $f(u_i^\sigma) = f(u_j^\sigma) = c$.

P2. $f(v) \leq c$ for all other vertices of $H_{Del_G}(O)$ not contained in the affine hull of $e_{i,j}$.

Let g be the linear function $g(x) = f(x) - f(O)$. By property P1, we have $f(u_j^\sigma - u_i^\sigma) = g(u_j^\sigma - u_i^\sigma) = 0$. This means that $f(b_{\sigma(j)} + \cdots + b_{\sigma(i+1)}) = g(b_{\sigma(j)}) + \cdots + g(b_{\sigma(i+1)}) = 0$. Either $g(b_{\sigma(j)}) = g(b_{\sigma(j-1)}) = \cdots = g(b_{\sigma(i+1)}) = 0$ or there exists $j \leq k_1, k_2 \leq i+1$ such that $g(b_{\sigma(k_1)}) > 0$ and $g(b_{\sigma(k_2)}) < 0$. In the first case, we have $f(u_{i+1}^\sigma) = f(u_i^\sigma) + g(b_{\sigma(i+1)}) = c$ but by property P2, this means that u_{i+1}^σ is contained in the affine hull of $e_{i,j}$ and hence $u_{i+1}^\sigma = \lambda \cdot (u_i^\sigma - u_j^\sigma)$ for some $\lambda \in \mathbb{R}$. But this contradicts the connectivity of G. In the second case, we have $f(u_{\sigma(i)} + b_{\sigma(k_1)}) = f(u_i^\sigma) + g(b_{\sigma(k_1)}) > c$ and since $u_{\sigma(i)} + b_{\sigma(k_1)}$ is a vertex of $H_{Del_G}(O)$, we obtain a contradiction. This concludes the proof. $\square$

Problem 4.1.13. *We have characterized the zero, one and $n-1$ dimensional faces of $H_{Del_G}(O)$. Can we obtain similar characterizations of the other faces of $H_{Del_G}(O)$?*

4.1.2 Combinatorics of $H_{Del_G}(O)$.

Recall that the f-vector of an n-dimensional polytope P is the vector $(f_0, \dots, f_{n-1})$ where f_k is the number of k-dimensional faces of P. In the previous section, we showed that $f_0(H_{Del_G}(O)) = 2^{n+1} - 2$ and $f_{n-1}(H_{Del_G}(O)) = n \cdot (n+1)$.

For each vertex v, we denote by $d_k(v)$, the number of k-dimensional faces incident to v.

Lemma 4.1.14. *For a vertex v of the form $\sum_{j=0}^{k-1} b_{i_j}$ where the indices i_js are all distinct, we have $d_{n-1}(v) = k \cdot (n+1-k)$.*

Proof. Consider a vertex v of the form $\sum_{j=0}^{k-1} b_{i_j}$ where i_js are all distinct and let $S_v = \{i_0, \dots, i_{k-1}\}$. The facets that contain v are of the form $F_{i,j}$ where $i \in S_v$ and $j \notin S_v$. There are $k \cdot (n+1-k)$ such facets. □

Corollary 4.1.15. *Let $g(k) = k \cdot (n+1-k)$, if $g(k_1) = g(k_2)$ for $k_1 \neq k_2$ then $k_1 + k_2 = n+1$.*

4.1.3 A Property of Affine Maps

Let $M(x) = A \cdot x + t$ for some non-singular linear transformation A and $t \in H_0$. Let $\mathcal{P}$ be an n-dimensional polytope in H_0. Then $M(\mathcal{P})$, the image of M on $\mathcal{P}$, is an n-dimensional polytope in H_0. Let $F_k(\mathcal{P})$ be the set of k dimensional faces of a polytope $\mathcal{P}$. We have the following property:

Lemma 4.1.16. *The map $M : F_k(\mathcal{P}) \to F_k(M(\mathcal{P}))$ is a bijection for k from 0 to n. Hence, if $\mathcal{F}$ is a k-dimensional face of $\mathcal{P}$ then $M(\mathcal{F})$ is also a k-dimensional face of $M(\mathcal{P})$.*

Corollary 4.1.17. *If v is a vertex of $\mathcal{P}$, then $M(v)$ is a vertex of $M(\mathcal{P})$ and $d_k(v) = d_k(M(v))$ for k from 0 to n.*

Recall that the polytopes $\mathcal{P}_1$ and $\mathcal{P}_2$ in H_0 are congruent if there exists an isometry M such that $\mathcal{P}_2 = M(\mathcal{P}_1)$. By an isometry, we mean a map of the form $M \cdot x = A \cdot x + t$ where A is an orthogonal transformation and $t \in H_0$

4.1.4 Correspondence between G and $H_{Del(G)}(O)$

For a convex polytope $\mathcal{P}$ in H_0, let $Aut(\mathcal{P})$ denote the automorphism group of $\mathcal{P}$. The following simple lemmas turn out to be useful.

Lemma 4.1.18. *Let $\mathcal{P}$ be a polytope in $\mathbb{R}^n$, let M be an element of $Aut(\mathcal{P})$, then M permutes the vertices of $\mathcal{P}$.*

Lemma 4.1.19. *Every element of $Aut(\triangle)$ is an orthogonal transformation, i.e., the translation part of the isometric map is zero.*

Proof. By Lemma 4.1.18, a transformation $M \in Aut(S)$ permutes the vertices of $\mathcal{P}$. Let $\sigma \in S_{n+1}$ be the permutation induced by M on the vertices set of S. We have $M(t_i) = t_{\sigma(i)}$. But then for all the vertices of S, we have:

$$t_i \cdot t_j = \begin{cases} n^2 + n, & \text{if } i = j, \\ -(n+1), & \text{if } i \neq j. \end{cases}$$

Since σ is permutation, $\sigma(i) = \sigma(j)$ if and only if $i = j$. Hence, $M(t_i) \cdot M(t_j) = t_i \cdot t_j$. Furthermore, since the set $\{t_0, ..., t_n\}$ spans H_0, it follows that $M(v) \cdot M(u) = v \cdot u$ for all v, u in H_0. Hence, M is an orthogonal transformation. □

Theorem 4.1.20. *Let G_1 and G_2 be undirected connected graphs. The polytopes $H_{Del_{G_1}}(O)$ and $H_{Del_{G_2}}(O)$ are congruent if and only if G_1 and G_2 are isomorphic.*

Proof. If G_1 and G_2 are isomorphic, then we know that there exists a permutation map or in other words $\sigma \in Aut(S)$ such that $Q(G_2) = \sigma \cdot Q(G_1) \cdot \sigma^{-1}$. This implies that we have $b_i^{G_2} = \sigma \cdot b_{\sigma^{-1}(i)}^{G_1}$. We now claim that $H_{Del_{G_2}}(O) = \sigma.H_{Del_{G_1}}(O)$. In order to see this observe that $\triangle_{\sigma_1}(G_2) = \sigma \cdot \triangle_{\sigma^{-1}\sigma_1}(G_1)$. By the definition of $H_{Del_G}(O)$, we have $H_{Del_{G_2}}(O) = \cup_{\sigma_1 \in S_{n+1}} \triangle_{\sigma_1}(G_2) = \cup_{\sigma_1 \in S_{n+1}} \sigma \triangle_{\sigma^{-1}\sigma_1}(G_1) = \sigma.H_{Del_{G_1}}(O)$. By Lemma 4.1.19, we know that σ is an orthogonal transformation and hence this implies that $H_{Del_{G_1}}(O)$ and $H_{Del_{G_2}}(O)$ are congruent.

Conversely, if $H_{Del_{G_1}}(O)$ and $H_{Del_{G_2}}(O)$ are congruent, then there exists an isometry $M(x) = A(x) + t$ for some orthogonal transformation A and $t \in H_0$ such that $H_{Del_{G_2}}(O) = M(H_{Del_{G_1}}(O))$. We know by Lemma 4.1.16 that M induces a bijection between the facets of $H_{Del_{G_1}}(O)$ and $H_{Del_{G_2}}(O)$. By Lemma 4.1.9 we know that a facet of $H_{Del_G}(O)$ is of the form $F_{i,j}^G$ for some $0 \leq i, j \leq n$. Consider an arbitrary facet $F_{i,j}^{G_2}$ of $H_{Del_{G_2}}(O)$ and let the facet $M(F_{i,j}^{G_2})$ of $H_{Del_{G_1}}(O)$ be $F_{i',j'}^{G_1}$. By Corollary 4.1.17, we know that M induces a bijection between the vertices of $F_{i,j}^{G_2}$ and the vertices of $F_{i',j'}^{G_1}$. By Corollary 4.1.17 we know that the $(n-1)$-th degree d_{n-1} is conserved and by Corollary 4.1.15 this means that either $b_i^{G_2} = M(\sum_{l=0}^n b_l^{G_1} - b_{i'}^{G_1}) = M(-b_{i'}^{G_1})$ or $b_i^{G_2} = M(b_{i'}^{G_1})$. In the first case, consider the map $M'(x) = -M(x)$. Indeed M' is an orthogonal transformation and since $H_{Del_{G_1}}(O)$ is a centrally symmetric polytope, we have $H_{Del_{G_2}}(O) = M(H_{Del_{G_1}}(O))$. Hence, we may assume without loss of generality that $b_i^{G_2} = M(b_{i'}^{G_1})$. By Corollary 4.1.15, this implies that $b_j^{G_2} = M(b_{j'}^{G_1})$. We know that for every edge e of facet $F_{i,j}^{G_2}$, $M(e)$ is an edge incident on $F_{i',j'}^{G_1}$. Moreover M induces a bijection between the edges of $F_{i,j}^{G_2}$ with b_i as a vertex and the edges of $F_{i',j'}^{G_1}$ with b'_i as a vertex. We know from Lemma 4.1.12 that the edges of $F_{i,j}^G$ incident on b_i^G are $(b_i^G, b_i^G + b_k^G)$ for $k \notin \{i, j\}$. With this information, we deduce that M induces a bijection between the vertices of the form $b_k^{G_2}$ where $k \notin \{i', j'\}$ and the vertices of the form $b_k^{G_1}$ where $k \notin \{i, j\}$. Hence, M induces a permutation σ between the vertices of $H_{Del_{G_2}}(O)$ that are of the form $b_j^{G_2}$ and the vertices of $H_{Del_{G_1}}(O)$ of the form $b_j^{G_1}$ for $0 \leq j \leq n$. Hence, we have $M(\sum_{k=0}^n b_k^{G_1}) = M(O) = \sum_{k=0}^n b_k^{G_2} = O$. Hence, $t = O$ and M is an orthogonal transformation and we have $M(b_i^{G_1}) \cdot M(b_j^{G_1}) = b_i^{G_1} \cdot b_j^{G_1}$. Putting these together, we have a permutation $\sigma \in S_{n+1}$ such that $b_i^{G_2} \cdot b_j^{G_2} = b_{\sigma(i)}^{G_1} \cdot b_{\sigma(j)}^{G_1}$ for all integers $0 \leq i, j \leq n$. This implies that $Q(G_2)Q^t(G_2) = Q(G_2)^2 = \sigma Q^2(G_1)\sigma^{-1} = (\sigma Q(G_1)\sigma^{-1})(\sigma Q(G_1)\sigma^{-1})$. But, we know that $Q(G_1)$ and $Q(G_2)$ are positive semidefinite and hence $\sigma Q(G_1)\sigma^{-1} = \sigma Q(G_1)\sigma^t$ is also positive semidefinite. By the unique squares lemma [47], we can conclude that $\sigma Q(G_1)\sigma^{-1}$ is the unique positive semidefinite square root of $Q^2(G_2)$ and hence $Q(G_2) = \sigma Q(G_1)\sigma^{-1}$. This shows that G_1 and G_2 are isomorphic. □

Remark 4.1.21. There are simpler constructions that also have the property shown in Theorem 4.1.20. For example for every connected graph G on $n+1$ vertices, associate an n-dimensional simplex given by $\mathcal{S}(G) = CH(b_0, \ldots, b_n)$ where $b_0, \ldots, b_n$ are the rows of the Laplacian of G. A argument similar to last part of the proof of Theorem 4.1.20 shows that: Let G_1 and G_2 be connected graphs $\mathcal{S}(G_1)$ is congruent to $\mathcal{S}(G_2)$ if and only if G_1 and G_2 are isomorphic. We will explore this direction in more detail in Chapter 6. On the other hand, Theorem 4.1.20 is more "canonical" in connection to the scarf complex and the polytopes $H_{Del(G)}(O)$ have a geometric interpertation in terms of the Laplacian lattice while the simplex $\mathcal{S}(G)$ does not seem to have any direct interpertation.

Remark 4.1.22. We know that the lengths of the edges of $H_{Del_G}(O)$ is essentially the degrees of different vertices of G. The volume of $H_{Del_G}(O)$ is essentially the number of spanning trees of $H_{Del_G}(O)$. Is there such an interpretation for the volumes of the other faces of $H_{Del_G}(O)$ in the appropriate measures?

4.2 On the Number of Graphs with a given Laplacian Lattice

In Lemma 4.1.2 of the previous section, we observed that the Laplacian lattice of any tree is the root lattice A_n. This observation raises the problem of counting the number of graphs that have A_n as their Laplacian lattice. The matrix-tree theorem gives an answer to the problem.

Lemma 4.2.1. *The number of graphs that have A_n as their Laplacian lattice is exactly $(n+1)^{n-1}$, i.e., the number of labelled trees on $n+1$ vertices.*

Proof. By Lemma 4.1.2 of the previous section, we know that the root lattice A_n is the Laplacian lattice of any tree on $n+1$ vertices. Conversely, any connected graph on $n+1$ vertices that is not a tree must contain at least two spanning trees and hence by Lemma 4.1.1 its Laplacian lattice must have covolume (with respect to A_n) strictly greater than one. □

Lemma 4.2.1 raises the following natural problem:

Given a sublattice L of A_n. Count the number of labelled connected graphs whose Laplacian lattice is L.

We denote the number of undirected connected graphs that have L as their Laplacian lattice as $N_{\mathrm{G}r}(L)$. Note that $N_{\mathrm{G}r}(L)$ is non-zero only if L is the Laplacian lattice of a connected graph. Our main result in this section is an upper bound on the number

$N_{\mathrm{Gr}}(L_G)$ for a Laplacian lattice L_G in terms of the number of different Delaunay triangulations of L_G under the simplicial distance function $d_\triangle$. As a corollary, we show that for a Laplacian lattice L_G of a multigraph with no zero entries in its Laplacian matrix, we have $N_{\mathrm{Gr}}(L_G) = 1$.

From now on, when we say Delaunay triangulation, we mean that the Delaunay triangulation under the simplicial distance function $d_\triangle$, see Subsection 3.1 for a discussion on Delaunay triangulation under polyhedral distance function and Subsection 3.2.3 for a discussion on the Delaunay triangulation of Laplacian lattices under the simplicial distance function $d_\triangle$.

We now make precise what "two triangulations are different" means.

For a triangulation T, let $H_T(O)$ be the union of simplices in the triangulation that have the origin O as a vertex. We say T_1 and T_2 are the same if $H_{T_1}(O) = H_{T_2}(O)$, otherwise they are different.

Theorem 4.2.2. *Let* $\mathrm{Del}(G)$ *be the Delaunay triangulation of* L_G *defined by the graph* G *under the distance function* $d_\triangle$ *and* $N_{\mathrm{Del}}(L_G, \triangle)$ *be the number of different Delaunay triangulations of* L_G*, we have* $N_{\mathrm{Gr}}(L_G) \leq N_{\mathrm{Del}}(L_G, \triangle)$.

4.2.1 Proof of Theorem 4.2.2

We know that every graph provides a Delaunay triangulation of its Laplacian lattice (Theorem 3.2.6). We show that if $H_{Del_{G_1}}(O) = H_{Del_{G_2}}(O)$ then the Laplacian matrices $Q(G_1)$ and $Q(G_2)$ of G_1 and G_2 are equal. In other words, G can be uniquely recovered from $H_{Del_G}(O)$. This would imply that two graphs cannot give rise to the same Delaunay triangulation and Theorem 4.2.2 then follows.

Lemma 4.2.3. *Let* $G_1,\ G_2$ *be connected graphs,* $H_{Del_{G_1}}(O) = H_{Del_{G_2}}(O)$ *if and only if* $Q(G_1) = Q(G_2)$.

Proof. Indeed if $Q(G_1) = Q(G_2)$ then $H_{Del_{G_1}}(O) = H_{Del_{G_2}}(O)$. To show the converse, we describe an algorithm to uniquely recover $Q(G_1)$ from $H_{Del_{G_1}}(O)$.

Define the set C_i as follows:

$C_i = \{p = (p_0, \dots, p_n) \in \mathbb{R}^{n+1} |\ p_i \geq 0 \text{ and } p_j \leq 0 \text{ for all } j \neq i\}$.

Indeed, we verify that C_i is a cone. Consider the set $H_{Del_{G_1}}(O)|C_i$ of vertices of $H_{Del_{G_1}}(O)$ that are contained in the cone C_i. Pick a vertex v, say in $H_{Del_{G_1}}(O)|C_i$ that maximizes the value of the i-th coordinate.

First, a vertex with this property exists since the vertex $b_i(G_1)$ is contained in C_i and $b_{ij}(G_1)$ the j-th coordinate of $b_i(G_1)$ satisfies $b_{ij}(G_1) \leq 0$ for $i \neq j$. Furthermore, we claim that v is unique and is equal to $b_i(G_1)$. Now, assume that $v \neq b_i(G_1)$. By Lemma 4.1.3, we know that $v = b_{i_0}(G_1) + \cdots + b_{i_k}(G_1)$ for some $0 \leq k \leq n-1$, i_js all being distinct. Denote $S_k = \{i_0, \dots, i_k\}$. Since $b_{ij}(G_1) \leq 0$ for all $i \neq j$ we know that $v_i = b_{ii}(G_1)$ and hence, $i \in S_k$ and $b_{ij}(G_1) = 0$ for all $j \in S_k \setminus \{i\}$. We know that $v_j \geq 0$ for all $j \in S_k$. But since $v \in C_i$, this means that $v_j = 0$ for all $j \in S_k \setminus \{i\}$. This means that $b_{jk}(G_1) = 0$ for all $j \in S_k \setminus \{i\}$ and $k \in (V(G) \setminus S_k) \cup \{i\}$. Hence, there are at least two components of G that are not connected, namely the induced subgraphs of vertex sets $S_k \setminus \{i\}$ and $(V(G) \setminus S_k) \cup \{i\}$. This contradicts our assumption that G is connected. This shows that if $H_{Del_{G_1}}(O) = H_{Del_{G_2}}(O)$, then $b_i(G_1) = b_i(G_2)$ for all $0 \leq i \leq n$ and hence $Q(G_1) = Q(G_2)$. This concludes the proof of the lemma.

$\square$

This shows that each graph G with $L_G = L$ contributes to a different Delaunay triangulation of L and hence, $N_{\mathrm{Gr}}(L_G) \leq N_{Del}(L_G, \triangle)$. This concludes the proof of Theorem 4.2.2. We know from Theorem 3.2.6 that a multigraph with no zero entries in its Laplacian matrix has a unique Delaunay triangulation. As a corollary we obtain:

Corollary 4.2.4. *If L_G is the Laplacian lattice of multigraph G such that every pair of vertices are connected by an edge, then $N_{\mathrm{Gr}}(L_G) = 1$.*

4.3 Covering and Packing problems

Covering and Packing problems on lattices have been widely studied, see Sloane and Conway [31] for a general introduction and Zong and Talbot [82] or Martinet [62] for a more specialised treatment of the subject.

Given a lattice L and a convex body $\mathcal{P}$ (typically this is a sphere), we study to how "optimally" do L-translates of C pack or cover $\mathbb{R}^{n+1}$. For a sublattice L of A_n, we define the packing density $\gamma_C(L)$ and the covering density $\theta_C(L)$ as:

$$\begin{aligned} \gamma_C(L) &= \mathrm{Pac}_C(L)/((n+1)\mathrm{Cvol})^{1/n}(L) \\ \theta_C(L) &= \mathrm{Cov}_C(L)/((n+1)\mathrm{Cvol})^{1/n}(L) \end{aligned} \tag{4.2}$$

where $\mathrm{Pac}_C(L)$ and $\mathrm{Cov}_C(L)$ is the packing and covering radius of L with respect to C and $\mathrm{Cvol}(L)$ is the covolume of the lattice with respect to A_n.

Remark that in the standard definition of packing and covering density, the volume of the lattice L appears in the place of $n+1$ times the covolume of L. Observe that the two notions are "equivalent" up to a factor that depends only on the rank of L and are interchangeable since we are interested in determining lattices with good packing and covering densities in a given rank and the order of the quotient group A_n/L is equal to the ratio of the volumes of A_n and L i.e. $\mathrm{Vol}(L) = \mathrm{Vol}(A_n).|A_n/L|$ (see Lecture V, Theorem 20 of Siegel [77] for a proof).

The lattice packing and covering problem respectively is to find lattices that pack $\mathbb{R}^n$ most densely and cover $\mathbb{R}^n$ most economically i.e., lattices that maximise $\gamma_C(.)$ and minimise $\theta_C(.)$. Explicit constructions of lattices that solve the sphere packing and covering problems are known only for lower dimensions [31].

We will use the fact that the covering and packing radius of the Laplacian lattice have a combinatorial interpretation in terms of the underlying graph to obtain a formula for the covering and packing density of the Laplacian lattice. We will use this information to show that the Laplacian lattices of graphs that are highly connected such as Ramanujan graphs have good packing and covering properties among Laplacian lattices of graphs.

Proposition 4.3.1. *The packing radius and covering radius of the Laplacian lattice L_G respectively are:*

$$\gamma_{\triangle}(L_G) = \mathrm{MC}_1(G)/((n+1)(\prod_{i=1}^{n} \lambda_i)^{1/n}) \tag{4.3}$$

$$\theta_{\triangle}(L_G) = (\sum_{i=1}^{n} \lambda_i)/(2(n+1)(\prod_{i=1}^{n} \lambda_i)^{1/n}) \tag{4.4}$$

where $\lambda_1 \leq \lambda_2 \cdots \leq \lambda_n$ are the non-zero eigenvalues of the Laplacian matrix of G.

This immediately gives a lower bound for the covering density of a Laplacian lattice

Theorem 4.3.2. (Lower bounds for the covering density of Laplacian lattices) *The covering density of the Laplacian lattice is at least $n/2(n+1)$.*

Proof. By Corollary 4.3.1, we know that $\theta_{\triangle}(L_G) = (\sum_{i=1}^{n} \lambda_i)/(2(n+1)(\prod_{i=1}^{n} \lambda_i)^{1/n})$ where $\lambda_1 \leq \lambda_2 \cdots \leq \lambda_n$ are the non-zero eigenvalues of the Laplacian matrix. Now we use the fact that the Laplacian matrix is a positive semidefinite matrix along with the AM-GM inequality to obtain that $\theta_{\triangle}(L_G) \geq n/2(n+1)$. □

We consider the problem of minimising the covering density and maximising the packing density of the Laplacian lattice over all connected graphs with a given number of vertices. First we consider the covering density case. Suppose that $\lambda_1, \dots, \lambda_n$ are arbitrary positive real numbers then the quantity $\mathrm{Cov}_{\triangle}(L_G)$ is minimised if $\lambda_1 = \lambda_2 = \cdots = \lambda_n$. But then in our case, these numbers are eigenvalues of the Laplacian matrix. Nevertheless, we would like the eigenvalues of the Laplacian matrix to be "clustered". This suggests that graphs with good expansion properties would be suitable. To make this intuition precise, we need the notion of a Ramanujan graph, see the survey of Horory et al. [48] for a more detailed discussion of the topic.

Definition 4.3.3. *A d-regular graph is called a Ramanujan graph if $\lambda^A(G) \leq 2\sqrt{d-1}$, where $\lambda^A(G) = max\{|\lambda_2^A|, |\lambda_{n+1}^A|\}$ and $d = \lambda_1^A \geq \cdots \geq \lambda_{n+1}^A$ are the eigenvalues of the adjacency matrix of G.*

Using the fact that $\lambda_i = d - \lambda_i^A$ we have:

Lemma 4.3.4. *The non-zero eigenvalues of the Laplacian matrix of a Ramanujan graph are located in the interval $[d - 2\sqrt{d-1}, d + 2\sqrt{d-1}]$.*

Suppose that a d-regular graph is a Ramanujan graph, then we know that the eigenvalues of its Laplacian matrix are concentrated around d (the degree of the graph) in an interval of width $2\sqrt{d-1}$. More precisely, we have: $d - 2\sqrt{d-1} \leq \lambda_i \leq d + 2\sqrt{d-1}$ for every non-zero eigenvalue of the Laplacian. Using this information, we will obtain an upper bound on the covering density of a Ramanujan graph.

Lemma 4.3.5. *Let G be a d-regular Ramanujan graph, we have $d + 2\sqrt{d-1} \geq (\prod_{i=1}^{n} \lambda_i)^{1/n}$ $\geq d - 2\sqrt{d-1}$*

Proof. Since G is a d-regular Ramanujan graph we have $d-2\sqrt{d-1} \leq \lambda_i \leq d+2\sqrt{d-1}$ for every i from $1 \dots n$. Using this information, we have $(d+2\sqrt{d-1}) \geq (\prod_{i=1}^{n} \lambda_i)^{1/n} \geq (d-2\sqrt{d-1})$. $\square$

Theorem 4.3.6. (Covering Density of Ramanujan graphs) *Let G be a d-regular Ramanujan graph then $\theta_{\triangle}(L_G) \leq (\frac{d}{4(d-2\sqrt{d-1})})$.*

Proof. We have $\sum_{i=1}^{n} \lambda_i = (n+1) \cdot d/2$ and by Lemma 4.3.5 we have $\prod_{i=1}^{n} \lambda_i \geq d - 2\sqrt{d-1}$. By Corollary 4.3.1, we know that $\theta_C(L_G) = (\sum_{i=1}^{n} \lambda_i)/(2(n+1)(\prod_{i=1}^{n} \lambda_i)^{1/n})$ where $\lambda_1 \leq \lambda_2 \dots \leq \lambda_n$ are the non-zero eigenvalues of the Laplacian matrix. Hence, we obtain $\theta_{\triangle}(L_G) \leq (\frac{d}{4(d-2\sqrt{d-1}}))$. $\square$

Remark 4.3.7. Note that we crucially use the fact that the eigenvalues of a Ramanujan graph are concentrated in a small interval. It is not clear if we can obtain upper bounds on the covering density for a general graph.

We now consider the problem of maximising the packing density of the Laplacian lattice. The formula for packing density presented in Theorem 4.3.1 suggests that graphs that maximise the packing density have high minimum cut and a relatively small number of spanning trees. We will provide a lower bound on the packing density of Ramanujan graphs.

Theorem 4.3.8. *Let G be a d-regular Ramanujan graph, then $\gamma_{\triangle}(G) \geq \frac{(d-2\sqrt{d-1})}{2(n+1)(d+2\sqrt{d-1})}$.*

Proof. By Corollary 4.3.1 we have $\gamma_{\triangle}(G) = \frac{\mathrm{MC}_1(G)}{(n+1)(\prod_{i=1}^{n} \lambda_i)^{1/n}}$. By Lemma 4.3.5 we have: $\prod_{i=1}^{n} \lambda_i \leq d+2\sqrt{d-1}$. We now obtain a lower bound on $\mathrm{MC}_1(G)$. We observe that the size of the cut S can be written as $u_S \cdot Q(G) \cdot u_S^t/2$ where is the indicator vector of S i.e.

$$u_{Si} = \begin{cases} 1, \text{ if the vertex with index } i \text{ is in } S, \\ -1, \text{ otherwise.} \end{cases}$$

We now observe that:

$$\min_{S \notin \{V, \emptyset\}} \frac{u_S \cdot Q(G) \cdot u_S^t}{2} \geq \min_{u \in \mathbb{S}^{n+1}} \frac{u \cdot Q(G) \cdot u^t/2}{2} \cdot \min_{S \notin \{V, \emptyset\}} u_S \cdot u_S \geq \min_{u \in \mathbb{S}^{n+1}} \frac{u \cdot Q(G) \cdot u^t}{2} = \lambda_n/2$$

Now since the graph is a Ramanujan graph, we know that $\lambda_n \geq d - 2\sqrt{d-1}$. Hence, $\gamma_{\triangle}(G) \geq \frac{d-2\sqrt{d-1}}{2(n+1)(d+2\sqrt{d-1})}$. $\square$

We do not know if the converse of Theorem 4.3.6 and Theorem 4.3.8 also holds. More precisely, suppose that the covering density of the Laplacian lattice of a graph is upper bounded by a suitably chosen constant c, then is it true that the graph is Ramanujan? Similarly, suppose that the Laplacian lattice of a d-regular graph has a high packing density then does the graph have "high" connectivity? Another natural question is that whether the lower bound on the packing density that we obtained for the Laplacian lattice of a Ramanujan graph is the best possible. Note that a trivial upper bound for $\gamma_{\triangle}(G)$ is d for a d-regular graph G.

Chapter 5

Algorithmic Aspects of Riemann-Roch

We now turn to algorithmic questions related to the Riemann-Roch theory. In this chapter, we construct algorithms for computing the rank with the main result being an algorithm for computing the rank of a divisor on a finite multigraph that runs in polynomial time when the number of vertices of the multigraph is fixed. In particular, we obtain an algorithm for computing the rank of a divisor on a finite graph that runs in time $O(2^{n \log n})$ where n is the number of vertices of the graph. The key ingredients are a new geometric interpretation of rank combined with algorithms from the geometry of numbers. We conclude the chapter by showing that computing the rank of a divisor on a general sublattice of A_n i.e., the corresponding decision problem is NP-hard.

5.1 Computing the Rank of a divisor on a finite Graph

5.1.1 A Simplification

We shall first observe that by using the Riemann-Roch theorem we can restrict our attention to divisors of degree between zero and $g-1$. Firstly, a divisor of negative degree must have rank minus one. Furthermore, by the Riemann-Roch formula we have:

Lemma 5.1.1. *If the degree of D is strictly greater than $2g-2$, then $r(D) = deg(D) - g$.*

Proof. Observe that if the degree of D is strictly greater than $2g-2$ then the rank of $K-D$ is -1 and apply the Riemann-Roch theorem. □

Furthermore, we can compute the rank of divisors of degree between g and $2g-2$ by computing the rank of $K-D$, a divisor that has degree between zero and $g-1$ and then applying the Riemann-Roch theorem. Hence, we consider the problem of computing the rank of a divisor of degree between zero and $g-1$. In fact, we consider the decision version of the problem i.e., we want to decide (efficiently) if $r(D) \leq k$ for every k between zero and $g-1$; observe that such a procedure combined with a binary search over the parameter k will compute the rank in time $O(\ln(g))$ times the running time of the procedure.

5.1.2 A First Attempt at Computing the Rank

Let us discuss a first attempt at computing the rank. We will compute rank directly from its definition. We will use the fact that there is a polynomial time algorithm for testing if $r(D) \geq 0$ due to the independent work of Dhar [32] and Tardos [80].

Algorithm 1. *1. Enumerate all effective divisors of degree at most the degree of the divisor D.*

2. Find an effective divisor E of smallest degree such that $r(D-E) = -1$ by using Dhar's algorithm.

Theorem 5.1.2. *The running time of Algorithm 1 is $O(2^{n \ln g})$.*

The running time of the Algorithm 1 is not polynomial in the size of the input even for a fixed number of vertices since the quantity $2^{n \ln g}$ is not polynomially bounded in the size of the input. There is general interest in obtaining an algorithm that runs in polynomial time for a fixed number of vertices and furthermore, in obtaining a singly exponential time algorithm i.e., an algorithm with running time $2^{O(n)}\text{poly}(\text{size}(G))$. We will now undertake a deeper study of rank to obtain an algorithm that runs in time polynomial in the size of the input provided that the number of vertices is fixed. More precisely, our algorithm has running time $2^{O(n \log n)} poly(\text{size}(G))$. An important ingredient is a geometric interpretation of rank that we shall obtain in the following section.

5.1.3 A Geometric Interpretation of Rank

We start with the following formula for rank from Chapter 1, first shown in Baker and Norine [12].

Theorem 5.1.3. *For any divisor D, we have:*

$$r(D) = \min_{\nu \in \text{Ext}(L_G)} deg^+(D - \nu) - 1 \tag{5.1}$$

where $deg^+(D) = \sum_{i:D_i>0} D_i$.

5.1.4 A Sketch of the Approach

Let us now briefly sketch our approach to computing the rank: We start with the formula to compute the rank and proceed as follows: we run over all the permutations $\pi \in S_{n+1}$ and for each permutation π suppose that we could compute $\min_{q \in L_G} deg^+(D - v_\pi + q)$

in time that is possibly exponential but only in n, then we would obtain an algorithm with running time $O(f(n)\text{poly}(\text{size}(G)))$ for some function f. But, how do we compute $\min_{q \in L_G} deg^+(D - v_\pi + q)$? One hope would be to reduce the problem to a closest vector problem on lattices or more generally to integer programming. Fortunately, the integer programming problem has an algorithm that runs in time that is exponential only in n (the rank of the lattice). Such an algorithm would run in time $O(2^{n \log n}\text{poly}(\text{size}(G)))$. This approach requires a better understanding of the deg^+ function that we now obtain.

Definition 5.1.4. (**Orthogonal projections onto** H_k) *For a point $P \in \mathbb{R}^{n+1}$ we denote by $\pi_k(P)$ the orthogonal projection of P onto the hyperplane H_k.*

For the sake of simplicity, we first consider the case where the divisor has degree $g-1$. In this case, we observe that $deg^+(D-\nu) = \frac{\ell_1(D-\nu)}{2}$, and that $\ell_1(D-\nu) = \ell_1(\pi_0(D) - \pi_0(\nu))$. We denote the set $\pi_0(\text{Ext}(L_G))$ the orthogonal projection of $\text{Ext}(L_G)$ onto the hyperplane H_0 by $\text{Crit}(L_G)$ and indeed, the orthogonal projections of $\text{Ext}(L_G)$ are the local maxima of the distance function $h_{\triangle, L_G}$ we refer to Chapter 1 for more details.

Corollary 5.1.5. *For any divisor D with $deg(D) = g-1$, we have:*

$$r(D) = \min_{c \in Crit(L_G)} \frac{\ell_1(\pi_0(D) - c)}{2} - 1. \tag{5.2}$$

Taking cue from Corollary 5.1.5, it is natural to ask if there is a similar "distance function" type interpretation for divisors of degree between zero and $g-1$. We will answer this question in the affirmative, the relevant distance function, actually a family of distance functions is the following:

Definition 5.1.6. (**Degree-Plus Distance**) Let k be a positive real number. For points P and Q in H_0, we define the degree-plus distance between P and Q as

$$d_k^+(P,Q) = \sup\Big\{r \mid \triangle(P,r) \cap \triangle(Q, r+k) = \emptyset\Big\}.$$

where for a point $R \in H_0$ and a real number $r > 0$, $\triangle(R,r) = r \cdot \triangle(O) + R$.

Note that though d_k^+ does not appear to be a distance function at first glance, we will actually show that it can be realised by a sequence of polyhedral distance functions

We will now note some basic properties of the function d_k^+.

Lemma 5.1.7. (**Translation Invariance**) *For any points P, Q, and T in H_0 and for any positive real numbers r_1 and r_2 we have $\triangle(P,r_1) \cap \triangle(Q,r_2) = \emptyset$ if and only if $\triangle(P+T,r_1) \cap \triangle(Q+T,r_2) = \emptyset$.*

Proof. Assume that $\triangle(P,r_1) \cap \triangle(Q,r_2) \neq \emptyset$ and consider a point $S \in \triangle(P,r_1) \cap \triangle(Q,r_2)$. Now, $S = r_1 \sum_{i=1}^{n+1} \alpha_i v_i + P = r_2 \sum_{i=1}^{n+1} \beta_i v_i + Q$ for some $\alpha_i \geq 0$, $\beta_i \geq 0$ and $\sum_i \alpha_i = \sum_i \beta_i = 1$. Now this implies that $S + T = r_1 \sum_{i=1}^{n+1} \alpha_i v_i + P + T = r_2 \sum_{i=1}^{n+1} \beta_i v_i + Q + T$. Hence, $S + T \in \triangle(P+T, r_1) \cap \triangle(Q+T, r_2) \neq \emptyset$. The converse follows by symmetry. $\square$

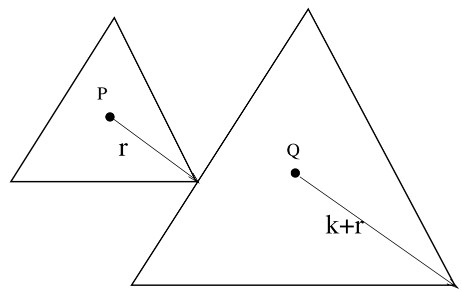

Figure 5.1: The distance function d_k^+

Lemma 5.1.8. (**Projection Lemma**) *Let R be a point in $\mathbb{R}^{n+1}$ with $deg(R) \geq 0$, let O be the origin and let $\pi_0(R)$ be the orthogonal projection of R onto H_0. We have*

$$\inf_{Z \in H^+(R) \cap H^+(O)} deg(Z) = (n+1) \sup \left\{ r \,|\, \triangle(\pi_0(R), r) \cap \triangle(O, r + \frac{deg(R)}{n+1}) = \emptyset \right\} + deg(R).$$

Proof. Consider a point X, say, in the intersection of $H^+(O)$ and $H^+(R)$ and consider the intersection of the hyperplane $H_{deg(X)}$ with $H^+(O)$ and $H^+(R)$. Observe that the intersection of $H_{deg(X)}$, the hyperplane $(1, \ldots, 1)^\perp$ translated by $deg(x)$ and $H^+(O)$ is a simplex that is a scaled and translated copy of $\triangle$, call it $\triangle_1$, centered at $\frac{deg(X)}{n+1}(1, \ldots, 1)$ and scaled by a factor of $\frac{deg(X)}{n+1}$. Similarly, the intersection of $H_{deg(X)}$ and $H^+(R)$ is also a simplex that is a scaled and translated copy of $\triangle$, call it $\triangle_2$ centered at $R + \frac{deg(X-R)}{n+1}(1, \ldots, 1)$ scaled by a factor of $\frac{deg(X-R)}{n+1}$. Observe that $deg(X) \geq deg(R) \geq deg(O)$. Indeed simplices $\triangle_1$ and $\triangle_2$ intersect at X and $deg(X)$ is equal to $n+1$ times the radius of $\triangle_2$ plus $deg(R)$. We now project the simplices $\triangle_1$ and $\triangle_2$ onto H_0 and obtain $\inf\{deg(Z) |\ Z \in H^+(R) \cap H^+(O)\} \geq (n+1) \sup\{r|\ \triangle(\pi_0(R), r) \cap \triangle(O, r + \frac{deg(R)}{n+1}) = \emptyset\} + deg(R)$. Now, consider a point P in $\triangle(\pi_0(R), r) \cap \triangle(O, r + \frac{deg(R)}{n+1})$ and observe that the point $X = P + (r + \frac{deg(R)}{n+1})(1, \ldots, 1)$ is a point in the intersection of $H^+(O)$ and $H^+(R)$. This shows that $\inf\{deg(Z) |\ Z \in H^+(R) \cap H^+(O)\} \leq (n+1) \sup\{r| \triangle(\pi_0(R), r) \cap \triangle(O, r + \frac{deg(R)}{n+1}) = \emptyset\} + deg(R)$. This completes the proof. $\square$

We are now ready to establish the connection between the deg^+ function and the function d_k^+.

Lemma 5.1.9. *For any pair of points P and Q in $\mathbb{R}^{n+1}$ with $deg(P) \geq deg(Q)$, we have*

$$deg^+(P - Q) = (n+1)\, d_k^+(\pi_0(P), \pi_0(Q)) + deg(P - Q)$$

for $k = \frac{deg(P-Q)}{n+1}$

Proof. First consider a point R in $\mathbb{R}^{n+1}$. We have $deg^+(R) = \sum_{R_i \geq 0} R_i = -deg(-R \oplus O)$, where $-R \oplus O = -(max(R_0, 0), \ldots, max(R_n, 0))$. Now we have:

$$-deg(-R \oplus O) = \inf_{Z \in H^+(R) \cap H^+(O)} deg(Z).$$

By Lemma 5.1.8 we have

$$\inf_{Z\in H^+(R)\cap H^+(O)} deg(Z) = (n+1)\sup\left\{ r \mid \triangle(\pi_0(R), r)\cap\triangle(O, r+\frac{deg(R)}{n+1}) = \emptyset\right\} + deg(R).$$

Now, for two points P and Q in $\mathbb{R}^{n+1}$, letting $R = P - Q$ in the above formula and applying Lemma 5.1.7, we obtain the relation given in the proposition. $\square$

The function d_k^+ is motivated naturally by the definition of the deg^+ function but is not very handy for geometric as well as computational reasons. In the following, we will obtain a more convenient representation of d_k^+. In fact, d_k^+ is closely related to the following family of polytopes: For a point $P \in H_0$ and $m, n > 0$, let $\mathcal{P}_{m,n}(P) = (\triangle(O, m) \oplus_{Mink} \bar{\triangle}(O, n)) + P$, where $\oplus_{Mink}$ denotes the Minkowski sum. Note that we use the notation $\oplus$ for the tropical sum.

Lemma 5.1.10. *For any positive real numbers m and n, $\mathcal{P}_{m,n}(P)$ is a convex polytope.*

Proof. Using the fact that Minkowski sum of two convex polytopes is a convex polytope and hence, $\mathcal{P}_{m,n}(O)$ is a convex polytope. Indeed translates of a convex polytope is also a convex polytope and hence, $\mathcal{P}_{m,n}(P)$ is also a convex polytope. $\square$

Lemma 5.1.11. *For points P and Q in H_0 and for $k \geq 0$, $d_k^+(P, Q) = \inf\{r| \ Q \in (\triangle(O, r) \oplus_{Mink} \bar{\triangle}(O, r + k)) + P\}$, where O is the origin.*

Proof. By definition $d_k^+(P, Q) = sup\{r| \ \triangle(P, r) \cap \bar{\triangle}(Q, r + k) = \emptyset\}$. Let $r_0 = d_k^+(P, Q)$ and consider a point R in the intersection of $\triangle(P, r_0)$ and $\bar{\triangle}(Q, r_0 + k)$. Rephrasing $d_{\triangle}(P, R) = r_0$ and $d_{\triangle}(Q, R) = d_{\bar{\triangle}}(R, Q) = r_0 + k$. This implies that $R - P \in \triangle(O, r_0)$ and $Q - R \in \bar{\triangle}(O, r_0 + k)$. This shows that $Q - P \in \triangle(O, r_0) \oplus_{Mink} \bar{\triangle}(O, r_0 + k)$ and we obtain $Q \in (\triangle(O, r_0) \oplus_{Mink} \bar{\triangle}(O, r_0 + k)) + P$. Hence, $d^+(P, Q) \geq \inf\{r| \ Q \in (\triangle(O, r) \oplus_{Mink} \bar{\triangle}(O, r + k)) + P\}$.

Furthermore, if Q is contained in $(\triangle(O, r) \oplus_{Mink} \bar{\triangle}(O, r + k)) + P$ then there exists a point $R = R_1 + R_2$ such that $R_1 \in \triangle(O, r)$ and $R_2 \in \bar{\triangle}(O, r + k)$ with $Q = R_1 + R_2 + P$ and we take R_3 such that $R_3 = Q - R_2 = R_1 + P$. Therefore, the point R_3 is contained in both $\triangle(Q, r + k)$ and $\triangle(P, r)$ and we obtain $\inf\{r| \ Q \in (\triangle(O, r) \oplus_{Mink} \bar{\triangle}(O, r + k)) + P\} \geq d_k^+(P, Q)$. This concludes the proof.

$\square$

As a corollary we obtain a handy characterization of the polytope $\mathcal{P}_{m,n}$:

Corollary 5.1.12. *Let P, Q be points in H_0, a point Q belongs to the polytope $\mathcal{P}_{r,r+d}(P)$ if and only if $deg^+(P + \frac{d(1,\ldots,1)}{n+1} - Q) \leq r(n + 1)$.*

Proof. Let a point Q be contained in $\mathcal{P}_{r,r+d}(P)$, then $\inf\{r'| \ Q \in (\triangle(O, r') \oplus_{Mink} \bar{\triangle}(O, r' + d)) + P\} \leq r$ and hence by Lemma 5.1.11 we know that $r \geq \inf\{r'| \ Q \in (\triangle(O, r') \oplus_{Mink} \bar{\triangle}(O, r' + d)) + P\} = d_d^+(P, Q)$. By Lemma 5.1.9, we know that $deg^+(P + \frac{d(1,\ldots,1)}{n+1} - Q) = (n + 1)d_d^+(P, Q)$. Hence, $deg^+(P + \frac{d(1,\ldots,1)}{n+1} - Q) \leq r(n + 1)$. Conversely, if a point Q satisfies $deg^+(P + \frac{d(1,\ldots,1)}{n+1} - Q) \leq r(n+1)$ then, $d_d^+(P, Q) \leq r$ and hence, $r \geq \inf\{r'| \ Q \in (\triangle(O, r') \oplus_{Mink} \bar{\triangle}(O, r' + d)) + P\}$ and hence $Q \in \mathcal{P}_{r,r+d}(P)$. $\square$

Now putting together, the formula for rank in Theorem 5.1.3, Lemma 5.1.9 and Lemma 5.1.11, we obtain the following geometric interpretation of rank:

Theorem 5.1.13. (**A Geometric Interpretation of rank**) *Consider a divisor D of degree d between zero and $g-1$, then D has rank $r_0 - 1$ if and only if $\pi_0(D)$ is contained in the boundary of the arrangement $\cup_{c \in \mathrm{Crit}(L_G)} \mathcal{P}_{r_1,r_2}(c)$ where $r_1 = r_0/(n+1)$ and $r_2 = (r_0 + g - 1 - d)/(n+1)$.*

Remark 5.1.14. The fact that $\mathcal{P}_{r_1,r_1+k}(c) \subseteq \mathcal{P}_{r_2,r_2+k}(c)$ if $r_1 \leq r_2$ is implicit in the statement of Theorem 5.1.13.

5.1.5 Computing the Rank for Divisors of Degree between zero and $g-1$

We now give an algorithm for computing the rank that runs in polynomial time for a fixed number of vertices. The algorithm uses two main ingredients:

1. The geometric interpretation of rank (Theorem 5.1.13).
2. Reduction to the algorithm for integer programming by Kannan [51].

For the sake of exposition, we first consider the slightly easier case of divisors with degree exactly $g-1$. We employ Theorem 5.1.5 to obtain the following algorithm:

Algorithm 2. *1. For each permutation $\pi \in S_{n+1}$, we compute $\min_{q \in L_G} \ell_1(\pi_0(D) - q)/2 - 1$ using Kannan's algorithm. The ℓ_1 unit ball is given to Kannan's algorithm as a separation oracle (we will provide an efficient implementation of the separation oracle in Lemma 5.1.15).*

2. We minimise over all permutations π.

We now turn to the general case: we start with the geometric interpretation for rank and we would like to reduce the problem to the integer programming problem. We construct a preliminary algorithm as follows:

Algorithm 3. *Find the smallest integer r such that $\pi_0(D)$ is contained in $\cup_{c \in \mathrm{Crit}(L_G)} \mathcal{P}_{r_1,r_2}(c)$ where $r_1 = \frac{r}{n+1}$, $r_2 = \frac{r+g-1+d}{n+1}$ by testing for all values of r from zero to $g-1$.*

Using the fact that the degree of the divisor is between zero and $g-1$, the algorithm would run in time $O(g \cdot 2^{O(n \log n)} \cdot poly(\mathrm{size}(G)))$. Since, g is not polynomially bounded in the size of the input, the algorithm does not run in polynomial time for fixed values of n. We resolve this problem by performing a binary search over the parameter r in the polytope $\mathcal{P}_{r_1,r_2}(c)$ and apply Kannan's algorithm at each step of the binary search. Since we know from Subsection 5.1.1, that the rank of the divisor is at most $g-1$ the algorithm terminates in $O(2^{n \log n} poly(\mathrm{size}(G)))$. Here is a formal description of the algorithm:

Algorithm 4. 1. *For each permutation $\pi \in S_{n+1}$, use binary search on the parameter r along with Kannan's algorithm to test if $\pi_0(D)$ is contained in $\cup_{q \in L_G} \mathcal{P}_{r,r+g-1-d}(c_\pi + q)$, the polytope is presented to Kannan's algorithm as a separation oracle.*

2. *Repeat over all permutations π.*

The straightforward way of presenting the polytope $\mathcal{P}_{r,r+k}$ to Kannan's algorithm is in terms of its facets. But since the number of facets of $\mathcal{P}_{r,r+k}$ is 2^{n+1}, the factor depending on n in the time complexity of the algorithm becomes larger than $2^{n \log n}$. Hence, we present the polytope $\mathcal{P}_{r,r+k}$ by a separation oracle to Kannan's algorithm and the following efficient implementation of the separation oracle ensures that the algorithm runs in time $2^{O(n \log n)} poly(\text{size}(G))$.

Lemma 5.1.15. (**A separation oracle for the polytope** $\mathcal{P}_{m,n}$) *There is a polynomial time separation oracle for the polytope $\mathcal{P}_{r,r+d}$ i.e., given any point p there is a polynomial time (in the bit length of p and the vertex description of $\mathcal{P}_{r,r+d}$) algorithm that either decides that p is contained in $\mathcal{P}_{r,r+d}$ or outputs a hyperplane separating the point p and the polytope $\mathcal{P}_{r,r+d}$.*

Proof. Given a point Q, compute the function $r' = deg^+(P + \frac{d(1,\dots,1)}{n+1} - Q)/(n+1)$ and if $r' \leq r$ then Q is contained in $\mathcal{P}_{r,r+d}(P)$ and otherwise let S^+ be the set of indices such that $P_i + \frac{d(1,\dots,1)}{n+1} > 0$, output the hyperplane H_S: $\sum_{x_i \in S^+}(-x_i + P_i + \frac{d(1,\dots,1)}{n+1}) \leq \frac{(r'+r)(n+1)}{2}$ as a separating hyperplane. To show that H_S is a separating hyperplane, assume the contrary and since $\sum_{x_i \in S^+}(-x_i + P_i + \frac{d(1,\dots,1)}{n+1}) = r'(n+1) > \frac{(r+r')(n+1)}{2}$ there is a point Q in $\mathcal{P}_{m,n}$ such that $\sum_{Q_i \in S^+}(P_i + \frac{d(1,\dots,1)}{n+1} - Q_i) \geq \frac{(n+1)(r'+r)}{2}$. We know that for the point $Q^L = Q - \frac{d(1,\dots,1)}{n+1}$ we have: $\sum_{j \in S^+}(Q^L)_j = r'(n+1) > \frac{(r+r')(n+1)}{2}$. Hence, $deg^+(Q^L) \geq \frac{(r+r')(n+1)}{2} > r(n+1)$ and $deg(Q^L) = d$, using Corollary 5.1.12 we obtain a contradiction.

□

The correctness of the algorithm is clear from Theorem 5.1.13.

Theorem 5.1.16. *For any divisor D with degree between zero to $g-1$, Algorithm 4 computes the rank of the divisor D.*

Theorem 5.1.17. *There is an algorithm (Algorithm 3) that computes the rank of a divisor on a finite multigraph G on n vertices with running time $2^{O(n \log n)} poly(size(G))$ and hence, runs in time polynomial in the size of the input for a fixed number of vertices.*

Proof. The first step in the algorithm takes time $O(\ln(g) 2^{O(n \log n)})\text{poly}(\text{size}(G))$ since a separation oracle for $\mathcal{P}_{m,n}$ can be constructed in polynomial time in the size of G and Kannan's algorithm takes $2^{O(n \log n)}\text{poly}(\text{size}(G))$ and we iterate $\ln(g)$ times. Since $\ln(g)$ is polynomially bounded in the size of the input, the time complexity of the algorithm is $O(2^{O(n \log n)}\text{poly}(\text{size}(G))$. □

5.1.6 The Algorithm

We now summarise the results that we obtained in the previous section to obtain an algorithm for computing the rank of a divisor.

Algorithm 5. *1. If $deg(D) < 0$ then, output $r(D) = -1$.*

2. If $g \leq deg(D) \leq 2g-2$ then, set $D' = K - D$ and compute $r(D')$ using 4. Output $r(D) = r(D') + deg(D) - (g-1)$.

3. If $deg(D) > 2g-2$, then $r(D) = deg(D) - g$.

4. If $0 \leq deg(D) \leq g-1$, then we invoke Algorithm 4 to compute $r(D)$.

Remark 5.1.18. In Chapter 1, we defined the notion of rank of a divisor for an arbitrary sublattice of the root lattice A_n. In such a general setting, we do not know if rank can computed in polynomial time even when the rank of the lattice is fixed. In our algorithm we crucially exploit our knowledge of the extremal points and the problem with handling the general case is that we do not have an explicit description of the extremal points as we have in the case of Laplacian lattices. As a consequence, we do not know how to find the extremal points in polynomial time even when the rank is fixed.

We will end this section by determining the vertices of the polytope $\mathcal{P}_{m,n}$.

Lemma 5.1.19. *The vertices of $\mathcal{P}_{R_1,R_2}$ are precisely of the form $w_{i,j} = R_1 t_i - R_2 t_j$ for $i \neq j$ where $t_0, \dots, t_n$ are the vertices of $\triangle$.*

Proof. Observe that $w_{i,j}$ is contained in $\mathcal{P}_{R_1,R_2}$ for all pairs i, j from 0 to n. We now show that $\mathcal{P}_{R_1,R_2}$ is contained in the convex hull of $w_{i,j}$ where i, j vary from 0 to n. Let p be a point in $\mathcal{P}_{R_1,R_2}$ by definition it can written as $R_1 \sum_{i=0}^{n} \lambda_i v_i - R_2 \sum_{i=0}^{n} \sigma_i v_i$ with $\lambda_i \geq 0$, $\sigma_i \geq 0$ for i from 0 to n and $\sum_{i=0}^{n} \lambda_i = \sum_{i=0}^{n} \sigma_i = 1$. We let $\ell_{ij} = \lambda_i.\sigma_j$ and write $p = \sum_{i=0}^{n} \sum_{j=0}^{n} \ell_{ij} w_{ij}$. We now verify that $\sum_{i,j} \ell_{i,j} = 1$ and $\ell_{i,j} \geq 0$. This shows that $\mathcal{P}_{R_1,R_2}$ is contained in the convex hull of $\{w_{i,j}\}_{i,j}$.

We now show that $w_{i,i}$ are not vertices of $\mathcal{P}_{R_1,R_2}$ since $w_{i,i}$ is contained in $\triangle(O, R_1 - R_2)$ and $\triangle(O, R_1 - R_2)$ is contained in $\triangle(O, R_1)$ if $R_1 \geq R_2$ and is contained in $\bar{\triangle}(O, R_2)$ otherwise. To conclude the proof of the lemma, it suffices to show that $w_{i,j}$ is a vertex if $i \neq j$. To this end, we consider the linear function $f_{i,j} = x_i - x_j$ and note that $w_{i,j}$ is the unique maximum of $f_{i,j}$ among all points in the set $\{w_{i,j}\}_{i,j}$. □

5.2 NP hardness results in the general case

Let L be a full-rank sublattice of A_n. We will now show that computing the rank function for general L is NP-hard. Actually we prove that deciding if $r(D) \geq 0$ is already NP-hard for general D and L.

By the results of Section 1.2, deciding if $r(D) = -1$ is equivalent to deciding whether $-D \in \Sigma(D)$ or not. So we will instead consider this membership problem. We will show below that this problem is equivalent to the problem of deciding whether a rational simplex contains an integral point. We then use this to show that it is generally NP-hard to decide if a given integral point D is contained in $\Sigma(L)$. As every point of positive degree is in $\Sigma(L)$, we may only consider the points of negative degree.

We first state the following simple lemma.

Lemma 5.2.1. *Let D be a point in $\mathbb{Z}^{n+1}$ of negative degree. We have $D \in \Sigma(L)$ if and only if the simplex $\bar{\triangle}_{\frac{-\deg(D)}{n+1}}(\pi_0(D))$ contains no lattice point (a point in L) in its interior.*

Proof. We saw in Section 1.4 that $\partial\Sigma^c(L)$ is the lower graph of the function $h_{\triangle,L}$. It follows that $D \in \Sigma(L)$ if and only if $-\frac{\deg(D)}{n+1} < h_{\triangle,L}(\pi_0(D))$. By the definition of $h_{\triangle,L}$, this means that $D \in \Sigma(L)$ is equivalent to $d_{\bar{\triangle}}(p, \pi_0(D)) = d_{\triangle}(\pi_0(D), p) > -\frac{\deg(D)}{n+1}$ for all $p \in L$, which is to say, $\bar{\triangle}_{\frac{-\deg(D)}{n+1}}(\pi_0(D))$ contains no lattice point. $\square$

Hence, the question of deciding whether $D \in \Sigma(L)$ is biols down to the following question:

Given a simplex of the form $\bar{\triangle}_r(x)$ with centre at x and radius $r \geq 0$, can we decide if there is a lattice point in the simplex?

A simple calculation shows that the vertices of $\bar{\triangle}_{-\frac{\deg(D)}{n+1}}(\pi_0(D))$ are all integral. This shows that with respect to the lattice L, the simplex $\bar{\triangle}_{-\frac{\deg(D)}{n+1}}(\pi_0(D))$ is rational, i.e., there exists a large integer N such that $N\bar{\triangle}_{-\frac{\deg(D)}{n+1}}(\pi_0(D))$ is a polytope with vertices all in L. (This is because L has full rank and itself integral.)

We now recall that the complexity of deciding if an arbitrary rational n-dimensional simplex in $\mathbb{R}^n$ contains a point of $\mathbb{Z}^n$ is NP-hard when the dimension n is not fixed, and it is polynomial time solvable when the dimension is fixed [15]. In our case, we are fixing the rational simplex, and L is an arbitrary sublattice of A_n. We now present a polynomial-time reduction from the problem of deciding if an arbitrary rational n-dimensional simplex in $\mathbb{R}^n$ contains a point of $\mathbb{Z}^n$ to the problem of deciding if $D \in \Sigma(L)$: Given the vertices $V(S) = \{v_1, \dots, v_n\}$ of a rational simplex S in $\mathbb{R}^n$, we do the following.

1. Compute the centroid $c(S) = \frac{\sum_i v_i}{n+1}$ of S and let $S' = S - c(S)$.
2. Define the linear map f from $\mathbb{R}^n$ to H_0 by sending $V(S')$ bijectively to $V(\bar{\triangle}) = \{e_0, \dots, e_n\}$. Let $\bar{\triangle}(x)$ be the image of S, where $x = f(c(S))$.
3. Let $L_0 = f(\mathbb{Z}^n)$ and N be a large integer such that $NL \subset A_n$ (such N exists since f and S are rational, and so L is rational). Remark that we have $NL \cap N\bar{\triangle}(x) \neq \emptyset$ if and only if $S \cap \mathbb{Z}^n \neq \emptyset$. Remark also that $N\bar{\triangle}(x) = \bar{\triangle}_N(Nx)$. Note that N can be choosen as the determinant of the matrix representing the simplex S and since the bitlength of the determinant is polynomially bounded in the total bitlength of the matrix $Nf(I)$ representing a basis of NL where I is the identity matrix and $Nf(I)$ is in turn polynomially bounded in the size of the input simplex S.

4. Let D be the integral point in $\mathbb{Z}^{n+1}$ defined by $D = Nx - N(n+1)(1, \ldots, 1)$. Then $\pi_0(D) = Nx$, $\deg(D) = -N(n+1)$, and $N\bar{\triangle}(x) = \bar{\triangle}_{-\frac{\deg(D)}{n+1}}(\pi_0(D))$.)

For L defined as above, we infer that $\bar{\triangle}_{-\frac{\deg(D)}{n+1}}(\pi_0(D)) \cap L \neq \emptyset$ if and only $S \cap \mathbb{Z}^d \neq \emptyset$. So we have

Theorem 5.2.2. *For an arbitrary full rank sublattice L of A_n, the problem of deciding if $r(D) = -1$ given a point $D \in \mathbb{Z}^{n+1}$ and a basis of L is* NP*-hard.*

As a consequence, we also note that the decision version of the problem of computing the rank is NP-hard.

Theorem 5.2.3. *Given an integer $k \geq -1$, a point $D \in \mathbb{Z}^{n+1}$ and a basis of L of a sublattice of A_n. The problem of deciding if $r(D) \geq k$ is* NP*-hard.*

It is interesting to note that for the case of Laplacian lattices of graphs on $n+1$ vertices, with a given basis formed by the n rows of the Laplacian matrix, the problem of deciding if an integral point belongs to the Sigma-Region can be done in polynomial time [45]. So we are naturally led to the following questions:

Question 5.2.4. *Given a full rank sublattice L of A_n, does there exist a special basis B of L such that if L is given with B, then the membership problem for the Sigma-Region of L can be solved in polynomial time ?*

Question 5.2.5. *Given a Laplacian sublattice of A_n, is it possible to find the special basis of L in time polynomial in n ? Given a sublattice of A_n, is it possible to decide if L is Laplacian in time polynomial in n?*

Chapter 6

Graph Automorphism

We start this chapter with a remark that we made in Chapter 4 that the Laplacian simplex of a graph, namely the convex hull of the rows of the Laplacian characterizes the graph completely up to isomorphism. Motivated by this observation, we study a dimensionality reduction approach towards the graph automorphism problem. Another result in this chapter is an exponential sum formula for counting automorphisms of a graph leading to a polynomial time algorithm that counts modulo p for any fixed prime p the number of automorphisms that violate a multiple of p-edges.

6.1 Dimensionality Reduction

6.1.1 A Simplex Associated with a Graph

Let $G = (V, E)$ be a connected undirected graph containing $n+1$ vertices. The Laplacian $Q(G)$ associated with G is the matrix $D - A$, where A is the adjacency matrix of G and D is the degree matrix, i.e., the diagonal matrix whose ith entry is the degree of the ith vertex. Recall that if G is a connected undirected graph on $n + 1$ vertices, then $Q(G)$ has rank n over $\mathbb{R}$.

Let $\triangle(G) \subseteq \mathbb{R}^{n+1}$ be the convex hull of the $n+1$ rows of $Q(G)$. We denote the vector $(1, \ldots, 1)$ by $\mathbf{1}$. We know that $\triangle(G)$ is contained in the hyperplane orthogonal to the vector $\mathbf{1}$. Indeed, $\triangle(G)$ is also a simplex.

Lemma 6.1.1. *If G is an undirected connected graph on $n+1$ vertices then $\triangle(G)$ forms an n-dimensional simplex whose centroid is at the origin.*

Using the fact that an isometric transformation f must induce a bijection between the vertices of the two simplices, we deduce that it must perserve centroids; since the centroid of every Laplacian simplex is the origin, we infer that the translation part of the isometric transformation f between two congruent Laplacian simplices is zero, i.e., f is just an orthogonal map.

Similarly, for the simplex $\triangle(G)$, that has its centroid at the origin, the set of all orthogonal transformations that map $\triangle(G)$ to itself forms a group, which we call the automorphism group of $\triangle(G)$, and each such transformation is called an automorphism.

Theorem 6.1.2. *Let G_1 and G_2 be undirected connected graphs on $n+1$ vertices. Then $\triangle(G_1)$ is congruent to $\triangle(G_2)$ if and only if G_1 and G_2 are isomorphic.*

Proof. ($\Leftarrow$) If G_1 and G_2 are isomorphic then there exists a permutation matrix σ of $n+1$ elements such that $Q(G_2) = \sigma Q(G_1)\sigma^t = \sigma(\sigma Q(G_1))^t$. Since σ is an orthogonal map, we conclude that σ is an isometric transformation that takes the Laplacian simplex of G_1 to the Laplacian simplex of G_2. Hence, $\triangle(G_1)$ and $\triangle(G_2)$ are congruent.

($\Rightarrow$) If $\triangle(G_1)$ and $\triangle(G_2)$ are congruent then there exists a permutation σ such that the Gram matrices of $Q(G_1)$ and $Q(G_2)$ satisfy $Q(G_2)Q(G_2)^t = \sigma Q(G_1)Q(G_1)^t\sigma^t$. Since $Q(G_1)$ and $Q(G_2)$ are symmetric matrices, we have

$$Q(G_2)^2 = (\sigma Q(G_1)\sigma^t)(\sigma Q(G_1)\sigma^t).$$

Now, observe that the LHS is a positive semi-definite symmetric matrix and that the matrices $Q(G_2)$ and $\sigma Q(G_1)\sigma^t$ are also positive semidefinite. Hence, we can apply the unique squares lemma ([47, p. 405]) to conclude that $Q(G_2) = \sigma Q(G_1)\sigma^t$. This shows that G_1 and G_2 are isomorphic as desired. $\square$

In fact, there is a family of simplices that have this property [72][1]. Note that the adjacency relationships of the graph is encoded in the Gram matrix as the inner-product. As a direct consequence of the theorem above we have the following:

Corollary 6.1.3. *The automorphism group of G is isomorphic to the automorphism group of $\triangle(G)$.*

Remark 6.1.4. Kaibel and Schwartz [50], and Akutsu [1] have also given geometric formulations of the graph isomorphism problem. The latter has shown that deciding congruence of point sets is at least as hard as graph isomorphism, by associating the following point set with a graph: each vertex is mapped to the one of the standard orthonormal vectors in $\mathbb{R}^n$, and the edge (v, w) is mapped to the midpoints of the vectors associated with v and w. For the more structured case of polytopes, Kaibel and Schwartz show that deciding congruence of polytopes is at least as hard as graph isomorphism by associating a simple n-polytope with a graph on $n+1$ vertices. Clearly, the size of the geometric object associated with the graph in both the cases is larger than the size of $\triangle(G)$ (especially in the construction of Kaibel and Schwartz, where the simple n-polytope can have $O(n^3)$ vertices in the worst case). Besides, certain geometric properties of $\triangle(G)$ have graph theoretic interpretations; as we mentioned before, using the Matrix-Tree Theorem (Theorem 13.2.1, [41]) we can show that the volume of $\triangle(G)$ is essentially (up to to a factor that depends on n) the number of spanning trees of G.

From now on in this chapter, we will focus on the problem of Graph Automorphism (GA), that is deciding whether a graph has a non-trivial automorphism. Clearly, Corollary 6.1.3 implies that checking GA is equivalent to checking the existence of an automorphism for the simplex associated with the graph. A *dimensionality reduction* approach

[1]The results here were derived independently of the results in [72].

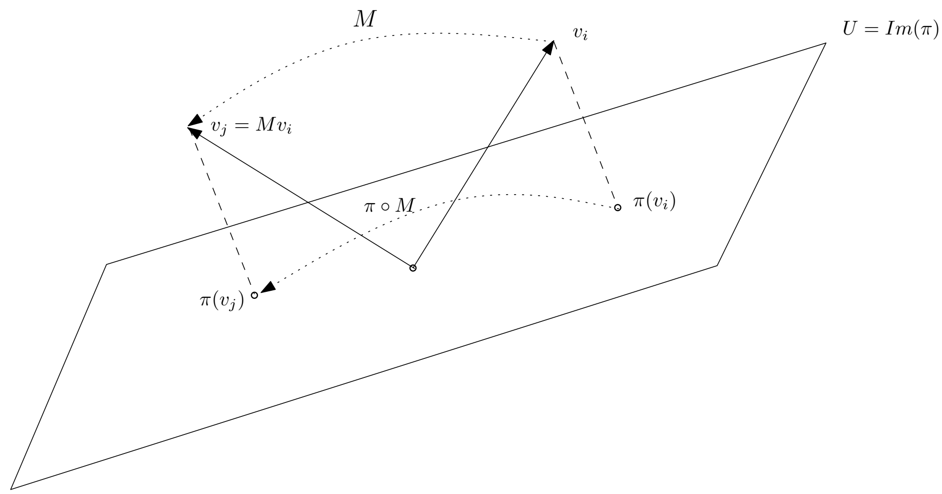

Figure 6.1: A projection π preserving an automorphism M.

to perform this check is the following: Project the simplex $\triangle(G)$ using a low distortion projection (such as those used in the Johnson-Lindenstrauss Lemma [DG03]) onto a space of smaller dimension, say $\mathbb{R}^k$; an automorphism of $\triangle(G)$ is transformed into an approximate automorphism (that is, each point is mapped to some ϵ-neighbourhood of another point) of projected simplex; we can then possibly enumerate all the approximate automorphisms in $\mathbb{R}^k$ and check if any of them yields an automorphism. The obstruction to this approach is that the number of approximate automorphisms are bounded by $O(n^{O(k^2)})$ [14], whereas the potential number of automorphisms of $\triangle(G)$ can be $n!$. Thus, a dimensionality reduction based approach has to necessarily drop automorphisms in the process. This observation, however, motivates the following question: *Given an automorphism M of $\triangle(G)$, what k-dimensional subspaces $U \subset \mathbf{1}^\perp$ of $\mathbb{R}^{n+1}$ "preserve" M?* More precisely, let $\pi : \mathbb{R}^{n+1} \to \mathbb{R}^{n+1}$ be the orthogonal projector operator with respect to U, and suppose M takes the point v_i of the simplex to the point v_j (i.e., $v_j = Mv_i$) then we say that U preserves M if

$$\pi(v_j) = \pi(M(\pi(v_i))), \tag{6.1}$$

as depicted in Figure 6.1. To answer our question, we will need the following notion: A subspace V is an invariant subspace of a linear map M if for all $w \in V$, $Mw \in V$. The theory of invariant subspaces is a rich one, and for more details we refer the reader to [42], [73]. We now briefly recall some standard results from the theory of invariant subspaces that turn to be useful for us and we include the proofs for the sake of reference.

The notion of U preserving M (see Equation 6.1) can be rewritten stated as

$$\pi(M(v_i)) - \pi(M(\pi(v_i))) = 0$$

or equivalently $\pi(M(v_i - \pi(v_i))) = 0$. Observe that $v_i - \pi(v_i) \in U^\perp$, and hence if $U^\perp$ is an *invariant subspace* of M then U preserves M. Conversely, if U preserves M then we know that for $i = 1, \ldots, n+1$, $M(v_i - \pi(v_i)) \in U^\perp$; but clearly, the $n+1$ vectors $v_i - \pi(v_i)$ generate $U^\perp$; thus, M fixes a basis for $U^\perp$, and hence $U^\perp$ is an invariant subspace of M. Hence, we have the following characterization:

Lemma 6.1.5. *A subspace U in $\mathbb{R}^{n+1}$ preserves M if and only if $U^{\perp}$ is an invariant subspace of M.*

For the special case of orthogonal transformations, we can replace $U^{\perp}$ by U in the result above, because of the following equivalence.

Lemma 6.1.6. *For an orthogonal transformation M, U is an invariant subspace of M if and only if $U^{\perp}$ is also an invariant subspace of M.*

Proof. Let $v_1, \ldots, v_k$ be a basis for the space U. Since M is an orthogonal map and U is an invariant subspace of M, we know that $Mv_1, \ldots, Mv_k$ forms a basis for U as well. For all vectors $w \in U^{\perp}$ we know that $\langle w, v_i \rangle = 0$, $i = 1, \ldots, k$. But as M is inner-product preserving we know that $\langle Mw, Mv_i \rangle = 0$ as well, that is, Mw is orthogonal to a basis of U, and hence $Mw \in U^{\perp}$, which implies that $U^{\perp}$ is an invariant subspace of M. The converse follows by replacing U by $U^{\perp}$. □

Another characterization of invariant subspaces of an orthogonal map is as follows.

Lemma 6.1.7. *Let π be the orthogonal projection operator for a subspace U. Then U is an invariant subspace for an orthogonal map M if and only if M and π commute with each other.*

Proof. Note that U is an invariant subspace of M if and only if $\pi M \pi = M\pi$. Let us first show that if π and M commute then U is an invariant subspace of M. Since $\pi M = M\pi$, we get $\pi M \pi = M\pi^2 = M\pi$ using the idempotent property of π. For the converse, we know that both U and $U^{\perp}$ are invariant subspaces of M, since M is orthogonal. Applying the result mentioned above to π and $I - \pi$, we get $M\pi = \pi M \pi = \pi M$. □

We have thus reduced the question of checking an automorphism of a graph G to finding invariant subspaces, possibly of low dimension, of an automorphism of $\triangle(G)$. This motivates our next section, where we study in more depth the invariant subspaces of an automorphism of $\triangle(G)$.

6.1.2 The Invariant Subspaces of an Automorphism of the Simplex

An automorphism M of $\triangle(G)$ induces a permutation σ on the vertex set of the simplex, such that if M takes v_i to v_j then σ maps i to j. The reduced form of this permutation gives rise to a family of invariant subspaces of M. To define these spaces, we need the notion of an orbit of an automorphism: An orbit ω of M is a subset $\{v_1, \ldots, v_k\}$ of the vertex set of $\triangle(G)$, such that $v_2 = Mv_1$, $v_3 = Mv_2, \ldots,$ $v_k = Mv_{k-1}$ and $v_1 = Mv_k$; thus the vectors in ω are those vectors whose indices come from an orbit in the reduced form of the permutation σ an equivalent way of writing ω is the set $\{v, Mv, \ldots, M^{k-1}v\}$. By the order of the orbit we mean the number of vectors contained in it.

For a set of vectors $S := \{v_1, \ldots, v_k\}$, let $\langle S \rangle$ or $\langle v_1, \ldots, v_k \rangle$ denote the linear space (over the reals) generated by $v_1, \ldots, v_k$. Then we have the following.

Lemma 6.1.8. *Let M be an automorphism of $\triangle(G)$.*

1. *If ω is an orbit of M then the subspace $\langle\omega\rangle$ is an invariant subspace of M.*

2. *Let $\omega_1, \ldots, \omega_m$ be the m orbits of M, and $\mathrm{ord}(M)$ the order of the cyclic subgroup generated by M when considered as an element of the automorphism group of $\triangle(G)$. Then $\mathtt{LCM}(|\omega_1|, \ldots, |\omega_m|) = \mathrm{ord}(M)$.*

3. *The invariant subspace spanned by an orbit ω of M contains an eigenvector corresponding to eigenvalue one. Let x_ω represent this eigenvector.*

4. *The invariant subspace spanned by an even order orbit ω of M contains an eigenvector corresponding to eigenvalue negative one. Let y_ω represent this eigenvector.*

5. *If a number ℓ divides the order of the orbit ω, then there exists an invariant subspace of dimension ℓ inside $\langle\omega\rangle$.*

6. *All the eigenvectors of M corresponding to eigenvalue one are contained in the space spanned by x_{ω_i}'s, for $i = 1, \ldots, m$. Similarly, all the eigenvectors of M corresponding to eigenvalue -1 are contained in the space spanned by y_{ω_i}'s, where ω_i's are the even order orbits of M.*

Proof. 1. The first result follows from the definition of the orbit. In the theory of invariant spaces, such subspaces are called *cyclic invariant subspaces* [42].

2. If $v \in \omega$, then we know that $M^{|\omega|}v = v$. Since the m orbits partition the vertex set we know that the smallest integer ℓ such that for all vertices v of the $\triangle(G)$, $M^\ell v = v$ is the $\mathtt{LCM}(|\omega_1|, \ldots, |\omega_m|)$. But ℓ is also equal to $\mathrm{ord}(M)$, by definition.

3. The vector $x_\omega := \sum_{v\in\omega} v$ is mapped to itself by M, and hence is an eigenvector of M corresponding to eigenvalue one.

4. Let $\omega = \{v, Mv, \ldots, M^{2k-2}v, M^{2k-1}v\}$. Then the vector $y_\omega :=(v - Mv) + (M^2v - M^3v) + \cdots + (M^{2k-2}v - M^{2k-1}v)$ is mapped to its negation by M.

5. Let $v_1, \ldots, v_k$ be the vertices in ω. Then we can partition the vertices into ℓ sets $S_0, S_1, \ldots, S_{\ell-1}$, where the vertex v_i goes into the set $S_{i \pmod{\ell}}$. Then it can be verified that the set
$$\{\sum_{v\in S_0} v, \sum_{v\in S_1} v, \ldots, \sum_{v\in S_{\ell-1}} v\}$$
generates a cyclic invariant subspace.

6. Let $v_1, \ldots, v_{n+1}$ be the vertices of $\triangle(G)$, and v be any vector in $\mathbf{1}^\perp$ such that $Mv = v$. Then since v_i's are affinely independent, we know that there is a unique representation of v in terms of v_i's, namely $v = \sum_{i=1}^{n+1} \alpha_i v_i$, where $\sum_i \alpha_i = 0$; the second property, in general, should be $\sum_i \alpha_i = 1$, but since the vertices satisfy

$\sum_i v_i = 0$, we can subtract a suitable scaling of this summation to get the property that $\sum_i \alpha_i = 0$. Let σ be the permutation induced by M. Then we have

$$Mv = \sum_i \alpha_i v_{\sigma(i)} = \sum_i \alpha_{\sigma^{-1}(i)} v_i.$$

Thus, $Mv = v$ implies that

$$\sum_i (\alpha_{\sigma^{-1}(i)} - \alpha_i) v_i = 0.$$

There are two possibilities: first, for all i, $\alpha_{\sigma^{-1}(i)} - \alpha_i = 0$, which implies that the α_i's are all the same for elements in a given orbit; second possibility is that $\alpha_{\sigma^{-1}(i)} - \alpha_i$ are all equal to some constant, but in this case we again claim that this constant is zero because $\sum_i \alpha_{\sigma^{-1}(i)} = \sum_i \alpha_i = 0$. Thus, v is a linear combination of the m eigenvectors x_{ω_i}, $i = 1, \ldots, m$. Note that the vectors x_{ω_i} are not linearly independent, since $\sum_{i=1}^{n+1} v_i = 0$. A similar argument holds the eigenvectors with eigenvalue -1 and the space spanned by y_{ω_i}'s; these vectors, however, are linearly independent.

$\square$

The invariant subspaces spanned by an orbit ω, though interesting, are very "local" in nature. That is, if we were to project $\triangle(G)$ orthogonally onto $\langle\omega\rangle$ and say G was mostly sparse, then most of the vertices not in ω will be mapped to the origin. As an alternative, one can use the closure property of invariant spaces under direct sums (see [42, p. 31]) to construct more "global" spaces; however, the dimension of these spaces is large; if in the worst case we were to take the direct sum of the invariant subspaces corresponding to all the orbits then we get an n-dimensional space. With dimensionality reduction as our aim we want to find some low dimensional invariant subspaces that are also "global" in nature. One such family of invariant subspaces is the Cross-Invariant Subspaces: Let $\omega_1, \ldots, \omega_m$ be the m orbits of an automorphism, v_i be an element in ω_i, and $a_1, \ldots, a_m$ be m real numbers. A cross-invariant subspace is the space generated by the vector $v := \sum_{i=1}^{m} a_i v_i$, i.e., the space $\langle v, Mv, M^2 v, \ldots\rangle$. Let us denote this space as $\mathcal{H}_a$, where $a \in \mathbb{R}^m$.

Clearly, cross-invariant subspaces will preserve an automorphism M of the simplex on projection. However, to recover the permutation corresponding to M from the permutation obtained by map induced by M on a cross-invariant subspace S, we need the subspace to be non-degenerate: A subspace $S \subseteq \mathbb{R}^{n+1}$ is said to be non-degenerate for a simplex $\triangle(G)$ if the orthogonal projection of the simplex onto S maps all the vertices to distinct points in S, and no point is mapped to the origin. Given the freedom in choosing a_i's in constructing cross-invariant subspaces, it is conceivable that non-degenerate cross-invariant subspaces exist; since if S is degenerate then a slight perturbation of S will restore non-degeneracy. The following theorem gives a more formal proof; the idea of the proof is similar to the proof technique used to show that certain family of hashing functions are universal (see [65] for more details).

Theorem 6.1.9. *An automorphism M of $\triangle(G)$ has non-degenerate cross-invariant subspaces $\mathcal{H}_a$, $a \in \mathbb{R}^m$.*

Proof. Let $\omega_1, \ldots, \omega_m$ be the orbits of M. Suppose $\mathcal{H}_a$ is the cyclic invariant subspace generated by the vector $w := \sum_{i=1}^m a_i w_i$, where w_i is some element in ω_i and the a_i's come from $[0, \ldots, N]$, for sufficiently large natural number N. A vertex $v \in \triangle(G)$ is mapped to the origin when projected into $\mathcal{H}_a$ if and only if it is orthogonal to the basis elements $w, Mw, M^2w \ldots$. We know that there is an index i such that v and w_i belong to the same orbit ω_i, i.e., there is a k such that $M^k w_i = v$.

Let us choose a_i's uniformly at random from $[0, \ldots, N]$. Then the probability that v is orthogonal to each of $w, Mw, \ldots,$ is smaller than the probability that v is orthogonal to $M^k w$, or equivalently that

$$a_k = -\frac{\sum_{j \neq k} a_j \langle v, w_j \rangle}{\|v\|^2}.$$

Since the a_i's are chosen uniformly at random from $[0, \ldots, N]$, the probability that a_k takes this special value is $1/(N+1)$. Thus, the probability that a vertex in the simplex is projected to the origin in $\mathcal{H}_a$ is smaller than $1/(N+1)$. We can similarly argue that the probability that two vertices $v, v' \in \triangle(G)$ are mapped to the same vertex in $\mathcal{H}_a$ (equivalently, that $v - v'$ is orthogonal to $\mathcal{H}_a$) is smaller than the probability that

$$a_k = -\frac{\sum_{j \neq k} a_j \langle v - v', w_j \rangle}{\langle v - v', v \rangle},$$

where k is such that $M^k w_i = v$; note that the denominator in the RHS is not zero because G is connected. Clearly, this probability is at most $1/(N+1)$.

Thus, the probability that $\mathcal{H}_a$ is non-degenerate is bounded by the sum of the probability that one of the points in $\triangle(G)$ is mapped to origin, and that a pair of points in $\triangle(G)$ is mapped to a single point. From the arguments above we know that this sum is smaller than

$$\frac{n+1+\binom{n+1}{2}}{N+1}.$$

Thus, by making N large enough we can ensure that in the space of cross-invariant subspaces there is a non-zero probability of picking a non-degenerate cross-invariant subspace $\mathcal{H}_a$. □

From now on, we use cross-invariant subspaces to implicitly imply that they are non-degenerate.

The results and concepts in this subsection, though developed only for orthogonal transformations, hold for general linear maps as well, as we did not invoke the orthogonality of the map M. In particular, we know that for a general permutation σ on the vertices of $\triangle(G)$, there exists a linear map (not necessarily orthogonal) that maps $\triangle(G)$ to itself. The results in this section also apply to this linear map. In the light of this similarity, we naturally ask the following question: What properties distinguish the invariant subspaces of an automorphism of $\triangle(G)$ from the invariant subspaces of a general linear map that maps $\triangle(G)$ to itself?

Given a permutation σ, let U_σ^- be the subspace spanned by the eigenvectors corresponding to eigenvalue -1, and U_σ^+ the space spanned by the eigenvectors corresponding to eigenvalue one. We start with characterizing order two permutations first.

Theorem 6.1.10. *Let σ be an order two permutation of G. Then σ is an automorphism of G if and only if U_σ^+ is orthogonal to U_σ^-.*

Proof. We first show that if σ is an automorphism of G then U_σ^+ is orthogonal to U_σ^-. The result follows if we show that for every element v in U_σ^-, and for $i = 1, \ldots, n+1$, $\langle v, v_i \rangle = -\langle v, v_{\sigma(i)} \rangle$. Let M_σ be the linear map associated with σ. Since M_σ is inner-product preserving, we know that

$$\langle v, v_i \rangle = \langle M_\sigma v, M_\sigma v_i \rangle = -\langle v, v_{\sigma(i)} \rangle.$$

For the converse, let $v_1, \ldots, v_k \in \Delta(G)$ be some elements from each of the k orbits of even order. For all $i = 1, \ldots, n+1$, define $w_{ji} := \langle v_j - v_{\sigma(j)}, v_i + v_{\sigma(i)} \rangle$. Since U_σ^+ and U_σ^- are orthogonal, $w_{ji} = 0$, for $j = 1, \ldots, k$. We claim that $w_{ji} = 0$ for all $j = 1, \ldots, n+1$; since $w_{ji} = -w_{\sigma(j)i}$, it follows that $w_{ji} = 0$ for all j appearing in the orbits of even order; for orbits of order one, w_{ji} is trivially zero, since $v_j = v_{\sigma(j)}$. Thus, $w_{ji} = 0$, for $i, j = 1, \ldots, n+1$. Summing the two equalities $w_{ij} = 0$ and $w_{ji} = 0$ we obtain that

$$\langle v_i, v_j \rangle = \langle v_{\sigma(i)}, v_{\sigma(j)} \rangle = \langle M_\sigma v_i, M_\sigma v_j \rangle$$

for all $i, j = 1, \ldots, n+1$. This implies that σ is inner-product preserving over a basis, we infer that M_σ must be an orthogonal map and hence σ is an automorphism of G. $\square$

A complete characterization of automorphisms in terms of their invariant subspaces is given in [9]: A permutation σ is an automorphism of the graph if and only if the eigenspaces of the adjacency matrix are invariant subspaces of σ. However, the key difference in their approach and ours is that in their approach eigenspaces are not guaranteed to be non-degenerate, whereas we are interested in finding non-degenerate invariant subspaces of the automorphism. A related question would be: Does projecting onto a slight perturbation of the eigenspace (to guarantee non-degeneracy) and then using some approximate congruence test (as in [4]) would suffice?

The Structure of the Invariant Subspaces Inside an Orbit Given a map M (not necessarily orthogonal) such that $M^k = I$, we want to study the space spanned by an orbit ω. As mentioned above, this space is a cyclic invariant subspace. An element in this space can be represented as the evaluation of a matrix polynomial $f(M)$, where $f(x) \in \mathbb{R}[x]$ is such that $\deg(f) < k$, evaluated at some point of the orbit. Clearly, this space is homomorphic to the ring $\mathbb{R}[x]/(x^k - 1)$. Moreover, the invariant subspaces contained inside ω are homomorphic to the rings $\mathbb{R}[x]/\langle g(x) \rangle$, where $g(x) \in \mathbb{R}[x]$ is a factor of the polynomial $x^k - 1$. We refer the reader to any standard algebra book to study the structure of this ring (for example [55]).

6.1.3 Applications

In this section, we describe two applications of the results obtained. In particular, we describe an algorithm that takes the Laplacian simplex of a graph and a cross-invariant subspace of an automorphism M as its input and finds an automorphism M' (not necessarily the same as M) in time $O(n^{O(\mathrm{ord}(M)^2)})$. As a corollary we have: given a cross-invariant subspace of an automorphism M of constant order, the algorithm constructs an automorphism M' in polynomial time.

Given $\triangle(G)$ and a cross-invariant subspace $\mathcal{H}_a$ for some automorphism of $\triangle(G)$, the algorithm proceeds as follows:

1. Project the vertices of $\triangle(G)$ orthogonally onto $\mathcal{H}_a$ to obtain a point set P.
2. Use the congruence algorithm in [4] to enumerate every linear map Q that takes P to itself.
3. For every such linear map Q test if the permutation induced by Q on P is an automorphism of $\triangle(G)$, until an automorphism M' is found.
4. Output M'.

The correctness of the algorithm is clear from the fact that the map induced by M on the subspace $\mathcal{H}_a$ induces an automorphism in $\triangle(G)$. Hence, the congruence algorithm finds an automorphism of $\triangle(G)$. The time complexity of the algorithm can be analysed as follows: The first step of the algorithm requires $n+1$ (the number of vertices of $\triangle(G)$) matrix-vector multiplications and hence the time taken can be upper bounded by $O(n^3)$. Denote the dimension of $\mathcal{H}_a$ by $d = O(\mathrm{ord}(M))$. The second and third steps of the algorithm together require $O(n^{2d^2+2})$ time, see [4]. Hence, the time complexity of the algorithm can be upper bounded by $O(n^{O(d^2)})$, or $O(n^{O(\mathrm{ord}(M)^2)})$.

Theorem 6.1.11. *Given a non-degenerate cross-invariant subspace of an automorphism M, it takes $O(n^{O(\mathrm{ord}(M)^2)})$ time to compute an automorphism of $\triangle(G)$ (not necessarily M).*

Remark 6.1.12. Using the Sylow theorem [55], we know that for a prime p, there is an automorphism of order p if and only if p divides the order of the automorphism group. Thus, if the order of the automorphism group is small then there are automorphisms that have small order, and hence invariant subspaces of small dimension. Constructing such invariant subspaces remains an open question. A possible approach for finding invariant subspaces of an order two automorphism could be based upon Theorem 6.1.10: since we know that projecting onto eigenvector with eigenvalue -1 preserves the automorphism, we can choose a random vector in the sphere in n-dimensions and perhaps it is close to one of the eigenvectors corresponding to some automorphism. This approach would have worked with high probability if the union of suitable ϵ-neighborhoods of such eigenvectors for all automorphisms was of measure at least half, however, this is not true (as the eigenvectors corresponding to non-automorphisms might be more denser) and so a purely random approach fails (even in the sense of getting approximate congruence).

We now discuss a second application of our results. We apply Theorem 6.1.2 to obtain a set of complete set of invariants that characterize a graph up to isomorphism. Consider the set of simplices in $\mathbb{R}^n$ and let T_n be the equivalence class of congruent simplices in $\mathbb{R}^n$. We know that T_n is a manifold and also a semi-algebraic set of dimension $n(n+1)/2$ [64]. We can use embedding theorems of manifolds, for example, Whitney's Embedding Theorem to infer that there exists an embedding of T_n into $\mathbb{R}^{n(n+1)+1}$. Thus, there exists a sequence of $n(n+1)$ real invariants that completely characterize a point in T_n and using Theorem 6.1.2 we have:

Theorem 6.1.13. *There exists a sequence of $n(n+1)+1$ invariants that uniquely characterize a connected undirected graph up to isomorphism.*

Remark 6.1.14. We finally note that it is possible to construct "trivial" invariants that uniquely characterize a graph up to isomorphism in the following way: since there are a finite number of equivalence classes of graphs on n vertices up to isomorphism, we assign a unique integer to each equivalence class. By construction these integers uniquely characterize a graph on n-vertices up to isomorphism. But these invariants do not vary continuously with respect to the manifold T_n, unlike the variants referred to in the theorem above.

A similar result was shown in [43], although the method used for constructing the invariants is different from our method.

6.2 Counting the Number of Automorphisms

In this section, we provide an exponential sum formulation for counting the number of automorphisms of a graph and show that the "constant order" terms of the exponential sum formulation can be computed in polynomial time. As an application of our result, we show that for a fixed prime p, we can count, modulo p, the number of permutations that violate a multiple of p edges in polynomial time. It is known that slightly more information such as the number of automorphisms modulo two is GA-hard, see [6].

Given a graph G, we construct a function on the set of permutations with the following property: the function vanishes for a permutation σ if and only if σ is an automorphism of G. The desired function is:

$$f(\sigma) = \frac{1}{2}\sum_{i=1}^{n+1}\sum_{j=1}^{n+1}(a_{i,j} - a_{\sigma(i),\sigma(j)})^2; \tag{6.2}$$

where $a_{i,j}$ is an entry in the adjacency matrix of G. It follows that $f(\sigma)$ has the following property:

Lemma 6.2.1. *If σ is an automorphism of G then the function $f(\sigma) = 0$; otherwise $f(\sigma)$ counts the number of edges violated by the permutation σ, i.e., the number of edges that are mapped to non-edges and vice-versa.*

Thus, we can interpret $f(\sigma)$ as an indicator function over the set of all permutations, S_n. Suppose $f(\sigma)$ was equal to some c for all non-automorphisms then the quantity $\sum_{\sigma \in S_n} f(\sigma)/c$ will give us the number of non-automorphisms, and the number of automorphisms can also be computed from this information. However, this assumption may not be true in general. To salvage this approach, we use the following standard property of exponential sums: For an integer m and a prime p we know that

$$\sum_{k=0}^{p-1} \exp(2\pi ikm/p) = \begin{cases} p & \text{if } p|m, \\ 0 & \text{otherwise.} \end{cases} \tag{6.3}$$

The proof is clear when $p|m$; otherwise, we observe that $\exp(2\pi ikm/p)$ are just the p roots of unity, and we know that their sum is zero. Using this property we have the following desired result.

Theorem 6.2.2. *For a sufficiently large prime p, the number of automorphisms N_A, is equal to*

$$\frac{1}{p} \sum_{\sigma \in S_n} \sum_{k=0}^{p-1} \exp(2\pi ikf(\sigma)/p). \tag{6.4}$$

We can choose p to be larger than $\max f(\sigma)$ *over all* $\sigma \in S_n$.

In number theory, the exponential sums studied are typically of the form

$$S := \sum_{x \in \mathbb{F}_p} \exp(2\pi if(x)/p),$$

where f is a polynomial of degree d and with coefficients in $\mathbb{F}_p$. The Weil character sum estimate states that under some mild technical conditions on the polynomial f, the summation S is upper bounded by $d\sqrt{p}$; note that a trivial upper bound for S is p. This result has applications in several other contexts, for example in derandomisation. See [8] for an introduction to exponential sums with some applications to computation. Though it seems like these exponential sums have the same flavour as the ones we consider, we are unaware of a more concrete connection between these two variants.

6.2.1 Computing the Exponential Sum

Let us assume that we can compute $\exp(\cdot)$ exactly. Then the straightforward approach to compute the sum in Equation (6.12) is to do the two summations. We will start by showing that the summation to p is polynomially bounded in n, or equivalently p is polynomially bounded. By Theorem 6.2.5, we require a prime p that satisfies the property that $p|f(\sigma)$ if and only if $f(\sigma) = 0$. This can be ensured if we choose p to be the smallest prime number greater than $\max_{\sigma \in S_n} f(\sigma)$. By Lemma 6.2.4, we know that $\max_{\sigma \in S_n} f(\sigma)$ can be upper bounded by $\frac{n(n-1)}{2}$. Moreover, by Bertrand's postulate [2], we know that there must be a prime between $\max f(\sigma) + 1$ and $2 \max f(\sigma)$. Hence, our choice of p is at most $n(n-1)$.

Now, let us look at the summation over S_n in Equation (6.12). If it is done naively then it will clearly take exponential time, as $|S_n| = n!$. In what follows, we will show that by interchanging the summations and expanding the exponential sum using the Taylor expansion we can compute lower order approximations to N_A efficiently. More precisely, using this approach we can rewrite N_A as

$$\begin{aligned} N_A &= \frac{1}{p}\sum_{k=0}^{p-1}\sum_{\sigma\in S_n}\sum_{\ell=0}^{\infty}\frac{1}{\ell!}\left(\frac{2\pi i k f(\sigma)}{p}\right)^{\ell} \\ &= \frac{1}{p}\sum_{\ell=0}^{\infty}\frac{1}{\ell!}\sum_{k=0}^{p-1}\sum_{\sigma\in S_n}\left(\frac{2\pi i k f(\sigma)}{p}\right)^{\ell}. \end{aligned}$$

Since we only need to consider the real part on the RHS, i.e., only the even values of ℓ, we obtain

$$\begin{aligned} N_A &= \frac{1}{p}\sum_{\ell=0}^{\infty}\frac{(-1)^{\ell}}{(2\ell)!}\sum_{k=0}^{p-1}\sum_{\sigma\in S_n}\left(\frac{2\pi k f(\sigma)}{p}\right)^{2\ell} \\ &= \frac{1}{p}\sum_{\ell=0}^{\infty}\frac{(-4\pi^2)^{\ell}}{p^{2\ell}(2\ell)!}\sum_{k=0}^{p-1}k^{2\ell}\sum_{\sigma\in S_n}f(\sigma)^{2\ell}. \end{aligned}$$

We now ask the following question: Till what value, L, of ℓ, do we need to expand the summation in ℓ to get a one-bit absolute approximation to N_A? We will show that $O(p^2 + n\log n)$ terms suffice. If L is such that the absolute value of the summation from L onwards is smaller than half then we are done, or if

$$\sum_{\ell=L}^{\infty}\frac{(4\pi^2)^{\ell}}{p^{2\ell}(2\ell)!}\sum_{k=0}^{p-1}k^{2\ell}\sum_{\sigma\in S_n}f(\sigma)^{2\ell} < \frac{1}{2}.$$

Since $p > f(\sigma)$ for all $\sigma \in S_n$, the above inequality will hold if

$$\sum_{\ell=L}^{\infty}\frac{(4\pi^2)^{\ell}}{(2\ell)!}\sum_{k=0}^{p-1}k^{2\ell}|S_n| < \frac{1}{2}.$$

Moreover, $\sum_{k=0}^{p-1} k^{2\ell} \le p^{2\ell+1}$. Thus, the above inequality follows if

$$\sum_{\ell=L}^{\infty}\frac{(4\pi^2)^{\ell}}{(2\ell)!}p^{2\ell+1}|S_n| < \frac{1}{2}. \tag{6.5}$$

Let us choose L large enough such that for all $\ell \ge L$

$$\frac{(4\pi^2)^{\ell}}{(2\ell)!}p^{2\ell+1}|S_n| \le 2^{-\ell},$$

because then we could use the geometric sum to obtain the Inequality ((6.5)). Using the fact that $2\ell! > \ell^\ell$, this follows if

$$L > \max(4\pi^2 p^2, \log(p|S_n|)), \tag{6.6}$$

or $L = O(p^2 + n\log n)$, since $|S_n| = n!$. Moreover, $p \leq n(n-1)$, thus L is polynomially bounded in n. This bound misleadingly suggests that it is possible to compute N_A in polynomial time. However, as we show next, this is not true in general. The hard part in computing N_A is in computing the summations

$$\Sigma_\ell := \sum_{\sigma\in S_n} f(\sigma)^\ell, \tag{6.7}$$

for ℓ even. In the rest of this section we give a method to compute Σ_ℓ in time $O(n^\ell)$.

Let us recall the definition of $f(\sigma)$ from Equation (6.10)

$$f(\sigma) = \frac{1}{2}\sum_{i=1}^{n+1}\sum_{j=1}^{n+1}(a_{i,j} - a_{\sigma(i),\sigma(j)})^2$$

Let

$$w_\sigma^{i,j} := (a_{i,j} - a_{\sigma(i),\sigma(j)})^2. \tag{6.8}$$

Then

$$\Sigma_\ell = \sum_{\sigma\in S_n}\left(\sum_{i=1}^{n+1}\sum_{j=1}^{n+1} w_\sigma^{i,j}\right)^\ell = \sum_{\sigma\in S_n}\left(\sum_{i,j}^{n+1} w_\sigma^{i,j}\right)^\ell.$$

Let W_σ be the $(n+1)^2$ dimensional vector of all $w_\sigma^{i,j}$'s. Using the multinomial theorem we obtain

$$\Sigma_\ell = \sum_{\sigma\in S_n}\sum_{I\in\mathbb{N}^{(n+1)^2}:|I|=\ell}\binom{\ell}{I}W_\sigma^I,$$

where $\binom{\ell}{I} := \ell!/(I_{1,1}!I_{1,2}!\cdots I_{n+1,n+1}!)$, and $|I|$ is the sum of all the entries in I. Since the inner-summation does not depend on σ, we can interchange the two summations to get

$$\Sigma_\ell = \sum_{I\in\mathbb{N}^{(n+1)^2}:|I|=\ell}\binom{\ell}{I}\sum_{\sigma\in S_n} W_\sigma^I. \tag{6.9}$$

For a given multi-index I, let S_n^I be the subset of those permutations σ, σ' in S_n such that $W_\sigma^I = W_{\sigma'}^I$. The following lemma gives us a more precise detailed characterization of S_n^I. For a multi-index I and a permutation σ, let us denote $\sigma(I)$ as the ℓ-dimensional vector I' that is obtained as follows: for each variable W_{ij}^σ appearing in the monomial W_σ^I we have $(\sigma(i),\sigma(j))$ as an entry in I'; since there are ℓ entries in I, the dimension of I' is 2ℓ.

Lemma 6.2.3. *For a given multi-index I, and σ, σ' in S_n, if $\sigma(I) = \sigma'(I)$ then $W_\sigma^I = W_{\sigma'}^I$.*

Proof. The condition $W^I_\sigma = W^I_{\sigma'}$ holds if for all entries in I we have $W^\sigma_{ij} = W^{\sigma'}_{ij}$. This would instead follow if $w^\sigma_{ji} = w^{\sigma'}_{ji}$. These two equalities follow if $\sigma(i) = \sigma'(i)$, $\sigma(j) = \sigma'(j)$. Thus for each of the ℓ non-zero entries in I, we have two corresponding values of i, j such that if σ and σ' are the same on these 2ℓ parameters then $W^I_\sigma = W^I_{\sigma'}$. □

We want to further simplify the term $\sum_{\sigma \in S_n} W^I_\sigma$. More precisely, we want to know, given a multi-index I, how many permutations σ can there be such that the monomial W^I_σ attains the same value? Note that given $I = [I_{i,j}]$, $i, j = 1, \ldots, n+1$, two permutations σ and σ' have the same weight W^I_σ, if and only if for all non-zero entries $I_{i,j}$ in I, $\sigma(i) = \sigma'(i)$ and $\sigma(j) = \sigma'(j)$. Equivalently, if $\mathcal{N}(I)$ is the set of distinct indices i, j appearing in the non-zero entries $I_{i,j}$ in I then both σ and σ' map $\mathcal{N}(I)$ to the same vector; note that the size of $\mathcal{N}(I)$ is at most 2ℓ as there can be at most ℓ non-zero entries $I_{i,j}$ and each can contribute two distinct values i, j to $\mathcal{N}(I)$. Let $J \in \{1, \ldots, n+1\}^{2\ell}$ and $N_{I,J}$ be the number of permutations that maps $\mathcal{N}(I)$ to J then we can rewrite

$$\sum_{\sigma \in S_n} W^I_\sigma = \sum_{J \in \{1,\ldots,n+1\}^{2\ell}} N_{I,J} \;\; W^I_{\mathcal{N}(I) \to J},$$

where $W^I_{\mathcal{N}(I) \to J}$ is a generic way of writing W^I_σ, for any permutation σ that maps $\mathcal{N}(I)$ to J.

Using this notation we can rewrite

$$\Sigma_\ell = \sum_{I \in \mathbb{N}^{\mathcal{N}} : |I| = \ell} \; \sum_{J \in \{1,\ldots,n+1\}^{2\ell}} N_{I,J} \;\; W^I_{\mathcal{N}(I) \to J}.$$

Clearly, the number of terms in the second summation is $(n+1)^{2\ell}$. The number of terms in the first summation is $(n+1)^{2\ell}$, because of the number of ways of choosing ℓ entries from a vector of dimension $(n+1)^2$. The crucial property of this reformulation of Σ_ℓ is that the number of terms appearing in it is $(n+1)^{4\ell}$, whereas the formulation in Equation (6.9) has $\Omega(n^n)$ terms. Thus, for fixed values of ℓ, we can compute Σ_ℓ efficiently, given $N_{I,J}$ can be computed efficiently. Given a graph G, we construct a function on the set of permutations with the following property: the function vanishes for a permutation σ if and only if σ is an automorphism of G. The desired function is:

$$f(\sigma) = \frac{1}{2} \sum_{i=1}^{n+1} \sum_{j=1}^{n+1} (a_{i,j} - a_{\sigma(i),\sigma(j)})^2; \tag{6.10}$$

where $a_{i,j}$ is an entry in the adjacency matrix of G. It follows that $f(\sigma)$ has the following property:

Lemma 6.2.4. *If σ is an automorphism of G then the function $f(\sigma) = 0$; otherwise $f(\sigma)$ counts the number of edges violated by the permutation σ, i.e., the number of edges that are mapped to non-edges and vice-versa.*

Thus, we can interpret $f(\sigma)$ as an indicator function over the set of all permutations, S_n. Suppose $f(\sigma)$ was equal to some c for all non-automorphisms then the quantity

$\sum_{\sigma \in S_n} f(\sigma)/c$ will give us the number of non-automorphisms, and the number of automorphisms can also be computed from this information. However, this assumption may not be true in general. To salvage this approach, we use the following property of exponential sums: For an integer m and a prime p we know that

$$\sum_{k=0}^{p-1} \exp(2\pi i k m/p) = \begin{cases} p & \text{if } p|m, \\ 0 & \text{otherwise.} \end{cases} \tag{6.11}$$

The proof is clear when $p|m$; otherwise, we observe that $\exp(2\pi i k m/p)$ are just the p roots of unity, and we know that their sum is zero. Using this property we have the following desired result.

Theorem 6.2.5. *For a sufficiently large prime p, the number of automorphisms N_A, is equal to*

$$\frac{1}{p} \sum_{\sigma \in S_n} \sum_{k=0}^{p-1} \exp(2\pi i k f(\sigma)/p). \tag{6.12}$$

We can choose p to be larger than $\max f(\sigma)$ *over all* $\sigma \in S_n$.

Computing $N_{I,J}$ and $W^I_{\mathcal{N}(I) \to J}$: We first remark that the computation of $N_{I,J}$ is independent of the graph and depends only on n and ℓ, and hence they can be precomputed based upon these parameters. More precisely, suppose we are given an $I \in \mathbb{N}^{(n+1)^2}$ and $J \subseteq \{1, \ldots, n+1\}^{2\ell}$ and we want to count the number of permutations σ in S_n that map $\mathcal{N}(I)$ to J. This latter equality imposes at most 2ℓ equalities on how σ behaves.

If there is a pair of equalities that is inconsistent, i.e., two equal elements map into different elements or two unequal elements map into the same element then set $N_{I,I'}$ to zero. If an element i maps to j and j does not map to i, we have an inconsistency and we set $N_{I,I'}$ to zero. Assume that the map I to I' is consistent then, the number of σ's that satisfy these equalities is precisely $(n+1-N_d(I))!$, where $N_d(I)$ is the number of distinct elements in I. We finally note that $(n+1-N_d(I))!$ can be computed in polynomial time.

If there is a pair of equalities that is inconsistent, i.e., violate the rules of a permutation then set $N_{I,J}$ to zero. In general, checking whether the map from $\mathcal{N}(I)$ to J comes from a permutation can be done in time polynomial in ℓ. Once this has been determined, it is straightforward to see that the number of permutations $N_{I,J}$ that satisfy the constraints imposed by $\mathcal{N}(I)$ and J is $(n+1-|\mathcal{N}(I)|)!$, which can be computed in polynomial time. Also, computing $W^I_{\mathcal{N}(I) \to J}$, a calculation that depends on the adjacency matrix of the graph, takes at most $O(\ell)$ steps. Thus, we have the following key result of this section.

Theorem 6.2.6. *Given $\ell \in \mathbb{N}$, the term $\sum_{\sigma \in S_n} f(\sigma)^\ell$, where the function f is defined in Equation (6.10) for a graph G, can be computed in time $O(n^{O(\ell)})\mathrm{poly}(\ell)$. In particular, for a constant ℓ, the time is polynomial in n.*

Based upon the theorem above, we can compute some interesting quantities.

Theorem 6.2.7. *Given a prime p, we can compute, modulo p, the number of permutations that violate a multiple of p edges in time $O(n^{O(p)})\mathrm{poly}(p)$. Hence, for a fixed prime p the computation takes time polynomial in the size of the input.*

Proof. By the construction of $f(\sigma)$ (see Equation ((6.10))), we know that $f(\sigma)$ counts the number of edges that the permutation σ violates. We compute the quantity Σ_{p-1}, which can be done in time $O(n^{O(p)})\mathrm{poly}(p)$. From Fermat's little theorem we know that $\Sigma_{p-1} \pmod p$ counts the number of permutations σ such that $f(\sigma)$ is relatively prime to p. Hence, $|S_n| - \Sigma_{p-1} \pmod p$ counts the number of permutations that violate a multiple of p edges. □

We note that a similar approach can be used to count the number, modulo p, of fixed-point free permutations that violate a multiple of p edges in polynomial time. The above computation seems to be on the border of what we can compute in polynomial time since we know that it is GA-hard to compute the number of automorphisms of a graph modulo two (see [6]), and it is $\oplus$P-hard to compute the number of fixed-point free automorphisms of order 2 of a graph modulo two, see [59, p. 16].

We conclude this section, by a brief discussion on computing the number of automorphisms modulo a special prime dependent on the graph.

Definition 6.2.8. *A prime number p is a good prime if for all non-automorphisms σ, p does not divide the number of edges violated by σ, i.e., p does not divide $f(\sigma)$.*

Using the ideas in the preceding theorem, we obtain the following:

Lemma 6.2.9. *For a good prime p, Σ_{p-1} is congruent modulo p to the number of non-automorphisms.*

We note that there exists good primes that are at most $2\max_{\sigma\in S_n} f(\sigma)$. Hence, there exist good primes that are polynomially bounded in n. But are there good primes upper bounded by a constant? It is unlikely that such small good primes exist since using the above approach we can also compute the number of fixed-point-free automorphisms of order two modulo a small prime p, but this problem is known to be $\#_k$P-complete for all $k \geq 2$ [59, Cor. 2,p. 17].

Chapter 7

Future Directions

We conclude by briefly discussing possible extensions of our work.

1. **Generalisation to higher dimensions:** The critical group of a graph is defined as the quotient group A_n/L_G where L_G is the Laplacian lattice of a graph. Recently [34], the notion of critical group has been generalised to higher-dimensions, i.e., to simplicial complexes, and higher-dimensional analogues of the Matrix-Tree theorem have been established. This progress raises the hope of generalising the Riemann-Roch theorem to simplicial complexes i.e., developing a discrete analogue of the Hirzebruch-Riemann-Roch theorem.

2. **Brill-Noether Theory for sublattices of A_n:** Let us now set-up a Brill-Noether theorem-type statement for sublattices of A_n. For integers r, d and g define, $\rho(g,r,d) = g - (r+1)(g-d+r)$. Is it true that if $\rho(g,r,d) \geq 0$, then for any full dimensional sublattice Λ of A_n with max-genus $g_{max} = g$, there exists an integral point of degree at most d and rank equal to r? Note that this part of Brill-Noether theorem for tropical curves is a direct corollary of Baker's Specialization Lemma, and the statement is true for Laplacian lattices of graphs as well [23]. We note that these proofs are not combinatorial, and use the classical Brill-Noether theorem; thus, they cannot be extended to more general lattices. (By using a particular type of tropical curves, a new proof of the non-existence result in the classical Brill-Noether theorem for algebraic curves (i.e., in the case $\rho(g,r,d) < 0$) was obtained in [29].)

   We believe that the geometric interpretation of rank obtained in Theorem 5.1.13 can be useful in answering this question for more general lattices.

3. **Extensions to Arithmetical graphs:** Recently, Lorenzini [58] considers Riemann-Roch structures on sub-lattices of rank n in $\mathbb{Z}^{n+1}$ which are perpendicular to a given vector $R = (r_0, \ldots, r_n)$ of $\mathbb{Z}^{n+1}$ with positive entries and with $\gcd\{r_i\}_{i=1}^{n} = 1$.

   He associates a genus $g(\Lambda)$ to any such lattice Λ, and considers the notions of canonical vector and Riemann-Roch structure for any such lattice.

As it is shown in [58, Section 5], there is a procedure that associates to any such lattice Λ, a sub-lattice Λ_0 of A_n, and the existence of canonical vectors and Riemann-Roch structure can be reduced to the corresponding questions about Λ_0.

For a sub-lattice Λ_0 of A_n, Lorenzini's g-number $g(\Lambda_0)$ coincides with our max-genus. In other words, it fairly easily follows that $g(\Lambda_0) = g_{max}(\Lambda_0)$. Lorenzini's definition of canonical vector is a relaxed version of our definition. Indeed, he only considers points of degree $g - 1$ ($=g_{max} - 1$) and call a point K of degree $2g - 2$ canonical if for any point D of degree $g - 1$, either both D and $K - D$ are equivalent to an effective point or neither is equivalent to an effective point. It is again fairly easy to show that in the case $g_{min} = g_{max}$, this definition coincides with our definition (but in general, the two definitions are different).

A nice class of examples of lattices Λ as above are given by arithmetical graphs. An arithmetical graph $\mathcal{A} = (G, M, R)$ is an undirected connected graph G on $n + 1$ vertices with the following additional data: M is an $(n + 1) \times (n + 1)$ equal to $D - A$ for a diagonal matrix D with strictly positive entries and A is the vertex-vertex adjacency matrix of the graph, $R \in \mathbb{N}^{n+1}$ is as before, so has strictly positive entries and $\gcd\{\mathrm{r_i}\} = 1$. In addition, R lies in the (right) kernel of M, i.e., $M \cdot R = 0$. An undirected graph is "naturally" an arithmetical graph with M being the Laplacian matrix of G and R being the vector $(1, \ldots, 1)$. Consider now the lattice $L_{\mathcal{A}}$ generated by the rows of the matrix M and let $L_{\mathcal{A},0}$ be the corresponding sub-lattice of A_n. It is easy to show that $L_{\mathcal{A},0}$ is the "Laplacian" lattice of a directed graph. In other words, there exists a directed multi-graph with "directed" adjacency matrix A_0 such that $L_{\mathcal{A},0}$ is generated by the rows of the corresponding Laplacian matrix $D_0 - A_0$, where D_0 is the diagonal matrix whose diagonal entries are given by the out-degrees of vertices. This leaves us with the following interesting question: is there a combinatorial interpretation of the extremal points of $L_{\mathcal{A},0}$ and more generally, is there a way to obtain a complete characterization of the extremal points of a sub-lattice of A_n defined by Laplacian of directed graphs? We note that, we gave a partial answer in this direction for Laplacian lattices of regular directed graphs, i.e., in the case where the out-degree of each vertex is equal to its in-indegree. Some partial results in this direction are obtained by Asadi and Backman [7].

4. **Zeta Function of graphs:** For an integer lattice Λ, as above, equipped with a Riemann-Roch structure, Lorenzini [58] also defines a zeta function and shows that it satisfies a functional equation similar to the zeta function on an algebraic curve. His construction gives rise to some natural questions about the correspondence between a lattice and its zeta function. For example, in an analogy with the case of projective curves, what can be said about two graphs which have the same zeta function? Examples of graphs with the same zeta-function but with different Jacobians are given in [58]. Understanding the relation between the Tutte polynomial of a graph and its zeta-function in further details is another interesting question.

   It is worth mentioning that there are other natural definitions of zeta functions for graphs and it is not clear at all how these different notions are related.

5. **Connections with other Riemann-Roch theorems and generalisations to convex polytopes:** Another intriguing general direction is to establish concrete connections between the Riemann-Roch theorems in other areas of mathematics. In particular, the Riemann-Roch theorem in toric geometry concerning counting lattice points in polytopes and Ehrhart polynomials has a similar flavour to Theorem 1. We believe that the key to making progress in this direction would be to obtain a better understanding of the function $r_{\mathcal{C}}(.,.)$, perhaps an interpretation in terms of counting lattice points in a polytope will be useful.

6. **Generalisation to polyhedral distance functions:** This direction is initiated by the observation that the information needed to describe our work is an n-dimensional lattice Λ and a regular simplex of dimension n that contains the origin in its interior. We now describe a reformulation of the Riemann-Roch theorem. Let Λ be a n-dimensional lattice and $\mathcal{C}(O,1)$ be a convex polytope of dimension n containing the origin, denoted by O. For a point P in the span of Λ and a positive integer r, we denote by $\mathcal{C}(P,r)$ the polytope $r \cdot \mathcal{C}(O,1) + P$ i.e., the copy of $\mathcal{C}(O,1)$ scaled by a factor of r and with the origin translated to P. Let $\bar{C}(P,r) = -r \cdot \mathcal{C}(O,1) + P$ and hence, if C is symmetric about the origin, then $\bar{C}(P,r) = C(P,r)$.

   For a general convex polytope $\mathcal{C}$, the distance function defined by $\mathcal{C}$ is $d_{\mathcal{C}}(P_1, P_2) = \inf\{t |\ P_1 \in \mathcal{C}(P_2, t)\}$. Consider the function $h_{\mathcal{C},\Lambda}(P) = \min_{Q \in \Lambda}\{d_{\mathcal{C}}(P,Q)\}$ and let $\mathrm{Crit}_{\mathcal{C}}(\Lambda)$ be the set of local maxima of the distance function $h_{\mathcal{C},\Lambda}(.)$. For a point $c \in \mathrm{Crit}_{\mathcal{C}}(\Lambda)$, the depth of c is the distance between c and a lattice point closest to c in the distance function $d_{\bar{\mathcal{C}}}$. For a real number $d > 0$ and a point P in the span of Λ, we define:

   $$r_{\mathcal{C}}(d,P) = \inf\{s |\ (\mathcal{C}(P,d) \oplus_{Mink} \bar{\mathcal{C}}(P,s)) \cap \mathrm{Crit}_{\mathcal{C}}(\Lambda) \neq \emptyset\}. \tag{7.1}$$

   Recall that the operation $\oplus_{Mink}$ is the Minkowski sum. The integer d essentially plays the role of degree of a configuration. By Theorem 5.1.13, we know that the definition of $r(k,D)$ reduces to the definition of rank if we choose $\mathcal{C}$ to be the regular simplex $\triangle$ and that Theorem 1 can be reformulated in terms of $r_{\mathcal{C}}(k,D)$ as follows:

   Theorem 7.0.10. *For a lattice Λ with rank n, if there is a point T in the span of Λ such that $\mathrm{Crit}_{\triangle}(\Lambda) = -\mathrm{Crit}_{\triangle}(\Lambda) + T$ and the points of $\mathrm{Crit}_{\triangle}(\Lambda)$ all have the same depth g_Λ then, for every point $P \in \mathbb{R}^n$ and for every positive integer k:*

   $$r_{\mathcal{C}}(k,P) - r_{\mathcal{C}}(2g_\Lambda - k, T - P) = k - g_\Lambda. \tag{7.2}$$

7. **Characterizing graphs with well distributed invariant subspaces:** Characterize graphs for which a randomly sampled subspace is an invariant subspace of an automorphism with high probability. In other words, characterize graphs for which the invariant subspaces of automorphisms are "well-distributed".

Appendix

In the first section of the appendix, we discuss a reformulation of Theorem 1.2.6 as a duality theorem for arrangement of simplices and in the second section we briefly describe facts that we exploited in Chapter 3 about the limiting behaviour of sequences of lattices.

A Duality Theorem for Arrangements of Simplices

Let L be a sublattice of A_n of rank n. For a real number $t \geq 0$, define the arrangement $\mathcal{A}_t$ as the union of all the simplices $\triangle_t(c)$ for $c \in \mathrm{Crit}(L)$, i.e.,

$$\mathcal{A}_t := \bigcup_{c \in \mathrm{Crit}(L_G)} \triangle_t(c).$$

A second arrangement $\mathcal{B}_t$ is defined as the union of all the simplices $\bar{\triangle}_t(p)$ for $p \in L$, i.e.,

$$B_t := \bigcup_{p \in L_G} \bar{\triangle}_t(p).$$

(Recall that $\bar{\triangle} = -\triangle$.)

Definition 7.0.11. *The covering radius of a lattice L denoted by $Cov(L)$ is the smallest real k such that $B_k = H_0$.*

Let G be an undirected graph on $n+1$ and with m edges (thus, $g = m - n$). Let L_G be the Laplacian lattice of G. (Recall that in this case, the covering radius of $\mathrm{Cov}(L_G)$ is the density of the graph.)

The two arrangements $\mathcal{A}$ and $\mathcal{B}$ are dual in the following sense.

Theorem 7.0.12. (Duality between $\mathcal{A}$ and $\mathcal{B}$) *For any $0 \leq t \leq \mathrm{Cov}(L_G)$, the arrangement $\mathcal{B}_t$ is the closure of the complement of the arrangement $\mathcal{A}_{\mathrm{Cov}(L_G)-t}$ in H_0, i.e.,*

$$\mathcal{B}_t = \Big(H_0 \setminus \mathcal{A}_{\mathrm{Cov}(L_G)-t}\Big)^c.$$

In particular, for any $0 \leq t \leq \mathrm{Cov}(L_G)$, $\partial\mathcal{B}_t = \partial\mathcal{A}_{\mathrm{Cov}(L_G)-t}$.

Proof. Let $x \in \mathcal{B}_t \cap \mathcal{A}_{\mathrm{Cov}(L_G)-t}$. By the definition of the two arrangements $\mathcal{B}$ and $\mathcal{A}$, there exists a point $p \in L_G$ and a point $c \in \mathrm{Crit}(L_G)$ such that $x \in \bar{\triangle}_t(p) \cap \triangle_{\mathrm{Cov}(L_G)-t}(c)$. By the triangle inequality for $d_\triangle$ and the results of Section 2.1, it follows that $d_\triangle(c,p) = \mathrm{Cov}(L_G)$, $d_\triangle(x,p) = t$, and $d_\triangle(c,x) = \mathrm{Cov}(L_G) - t$. Thus, we have $x \in \partial\mathcal{B}_t \cap \partial\mathcal{A}_{\mathrm{Cov}(L_G)-t}$. It follows that $\mathcal{B}_t$ and $\mathcal{A}_{\mathrm{Cov}(L_G)-t}$ have disjoint interiors, and so $\mathcal{B}_t \subseteq \Big(H_0 \setminus \mathcal{A}_{\mathrm{Cov}(L_G)-t}\Big)^c$. The other inclusion $H_0 \setminus \mathcal{A}_{\mathrm{Cov}(L_G)-t} \subset \mathcal{B}_t$ follows from the structural theorem of the Sigma-Region, Theorem 1.2.6 (and Theorem 1.2.7). Namely, we claim that for every point $x \in H_0$, there exists a point $p \in L_G$ and a point $c \in \mathrm{Crit}(L_G)$, such that $d_\triangle(c,x) + d_\triangle(x,p) = d_\triangle(c,p) = \mathrm{Cov}(L_G)$, and this clearly implies the inclusion $H_0 \setminus A_{\mathrm{Cov}(L_G)-t} \subset B_t$. Let p be a point of L_G such that $h_\triangle(x) = d_\triangle(x,p)$. By Proposition 1.4.12, the point $x - h_\triangle(x)(1,\ldots,1)$ lies on the boundary of Σ^c. By Theorem 1.2.7, there exists an extremal point ν of $\mathrm{Ext}^c(L_G)$ such that $\nu \le x - h_\triangle(x)(1,\ldots,1)$. Let c be the critical point $\pi_0(\nu) \in \mathrm{Crit}(L_G)$. Note that $h_\triangle(c) = \mathrm{Cov}(L_G)$. By Proposition 1.4.12, we have $\nu = c - \mathrm{Cov}(L_G)(1,\ldots,1)$. Thus, we have $c - h_\triangle(c)(1,\ldots,1) \le x - h_\triangle(x)(1,\ldots,1)$, or equivalently $c - (\mathrm{Cov}(L_G) - h_\triangle(x))(1,\ldots,1) \le x$. By the explicit definition of $d_\triangle$, we have $d_\triangle(c,x) \le \mathrm{Cov}(L_G) - h_\triangle(x) = \mathrm{Cov}(L_G) - d_\triangle(x,p)$. Since $d(c,p) \ge \mathrm{Cov}(L_G)$, this shows that $d_\triangle(c,x) = \mathrm{Cov}(L_G) - d_\triangle(x,p)$ and the claim follows. $\square$

Sequences of Lattices and their Limit

A sequence of lattices $\{L_n\}$ is said to **converge** to a lattice L_ℓ if for every $\delta > 0$ and $q \in L_\ell$ there exists a positive integer $N(\delta,q)$ and a family of bijective maps $\phi_N : L_\ell \to L_N$ for $N \ge N(\delta,q)$ such that $||q - \phi_N(q)||_2 < \delta$ for all $N \ge N(\delta,q)$. Similarly, a sequence of lattices $\{L_n\}$ is said to **uniformly converge** to a limit L_ℓ if for every $\delta > 0$ there exists an integer $N(\delta)$ and a family of bijective maps $\phi_N : L_\ell \to L_N$ for $N \ge N(\delta)$ such that $||q - \phi_N(q)||_2 < \delta$ for all $q \in L_\ell$ and for all $N \ge N(\delta)$. Note that since convex polytopes are topologically equivalent to Euclidean balls, the notion of convergence of a sequence of lattices with respect to polyhedral distance functions is equivalent to the above notion.

Typically, we take a basis $\mathcal{B}$ of the lattice and add an infinitesimal perturbation $\mathcal{B}^\epsilon$ to it and consider the lattice generated by the perturbed lattice. As ϵ tends to zero, the sequence of lattices converges to the lattice generated by $\mathcal{B}$, but not necessarily uniformly. Observe that there is a natural bijection ϕ_ϵ from L_ℓ to L_ϵ induced by the basis $\mathcal{B}$ defined as $\phi_\epsilon(\mathcal{B}\alpha) = \mathcal{B}^\epsilon\alpha$. In the following lemma we prove that the shortest vector and packing radius are preserved under the limit, see the book of Gruber and Lekkerkerker [44] for a more detailed treatment of lattices under perturbation.

Lemma 7.0.13. *Consider a sequence of lattices $\{L_n\}$ that converge to a lattice L_ℓ. For any convex polytope $\mathcal{P}$, the length of the shortest vector and packing radius under the distance function $d_\mathcal{P}$ is preserved under limit. More precisely we have:*

1. $\lim_{n\to\infty} \nu_\mathcal{P}(L_n) = \nu_\mathcal{P}(L_\ell)$
2. $\lim_{n\to\infty} \mathrm{Pac}_\mathcal{P}(L_n) = \mathrm{Pac}_\mathcal{P}(L_\ell)$.

Proof. i. Given an $\epsilon > 0$, we show that there exists a positive integer $n(\epsilon)$ such that $|\nu_{\mathcal{P}}(L_\ell) - \nu_{\mathcal{P}}(L_n)| \leq \epsilon$ for all $n \geq n(\epsilon)$. Consider the set of shortest vectors M_1 of the lattice L_ℓ. Let d be the the difference between the norm of the shortest vector and second shortest vector in the distance function $d_{\mathcal{P}}$. Let δ be the minimum of $d/2$ and $\epsilon/2$. Take a ball of radius R much larger than $\nu_{\mathcal{P}}(L_\ell)$ centered at the origin, in fact taking R to be $2\nu_{\mathcal{P}}(L_\ell) + d$ suffices. Using our assumption that $\{L_n\}$ converges to L_ℓ and the fact that number of points in L_ℓ that are contained inside the ball of radius R centered at the origin is finite, we know that there exists a positive integer $n(\delta)$ and a bijection $\phi_n : L_\ell \to L_n$ such that for $n \geq n(\delta)$ we have $d_{\mathcal{P}}(q, \phi_n(q)) < \delta$ for points q in L_ℓ that are contained in the ball of radius R. For a point p, we call $d_Q(O, p)$ the norm of the point. Now we claim that for all $n \geq n(\delta)$ any shortest vector of L_n belongs to $\mathcal{P}(q', \delta)$ where $q' \in M_1$. The argument is as follows: any element in L_n for $n \geq n(\delta)$ that does not belong to $\mathcal{P}(q', \delta)$ for $q' \in M_1$ must have norm strictly greater than $\nu_{\mathcal{P}}(L_\ell) + d/2$. On the other hand, we know that there exists an element of L_n that is contained in $\mathcal{P}(q', \delta)$ for some $q' \in M_1$ and hence has norm at most $\nu_{\mathcal{P}}(L_\ell) + \delta$. Since $\delta \leq d/2$, we arrive at a contradiction. We finally note that every point contained in the ball $\mathcal{P}(q', \delta)$ for $q' \in M_1$ has norm between $\nu_{\mathcal{P}}(L_\ell) + \delta$ and $\nu_{\mathcal{P}}(L_\ell) - \delta$. Hence, $|\nu_{\mathcal{P}}(L_n) - \nu_{\mathcal{P}}(L_\ell| \leq 2\delta \leq \epsilon$ for all $n \geq n(\delta)$.

ii. For the packing radius, the argument is essentially the same as the argument for shortest vectors except that in this case we carry out the argument on the $\mathcal{P}$-midpoints of the lattice points with the origin rather than on the lattice points. We consider the $\mathcal{P}$-midpoints of every point of L_ℓ with the origin. Let M_1 be the set of $\mathcal{P}$-midpoints of every point in L_ℓ that define its packing radius. Let d be the difference between the packing radius and the norm of $\mathcal{P}$-midpoints that are closer to the origin expect to the $\mathcal{P}$-midpoints that define the packing radius. Take δ to be the minimum of $d/2$ and $\epsilon/2$. Take a ball of radius R much larger than $\mathrm{Pac}_{\mathcal{P}}(L_\ell)$ centered at the origin. Since $\{L_n\}$ converges to L_ℓ and since the $\mathcal{P}$-midpoints vary continuously with perturbation, we know that there exists a integer $n(\delta)$ and a bijection $\phi_n : L_\ell \to L_n$ such that $d_{\mathcal{P}}(b, \phi_N(b)) < \delta$ for $\mathcal{P}$-midpoints in the ball of radius R. Now we claim that for $n \geq n(\delta)$ the $\mathcal{P}$-midpoint that defines the packing radius of L_n must be contained in $\mathcal{P}(b, \delta)$ for some $b \in M_1$. This follows from the fact that for $n \geq n(\delta)$ any element in L_n that does not belong to $\mathcal{P}(b', \delta)$ where $b' \in M_1$ must have norm strictly greater than $\mathrm{Pac}_{\mathcal{P}}(L_\ell) + d/2$. On the other hand, we know that the bisectors of elements of L_n that are contained in $\mathcal{P}(q', \delta)$ have norm at most $\mathrm{Pac}_{\mathcal{P}}(L_\ell) + \delta$. Since $\delta \leq d/2$, we arrive at a contradiction. We finally note that every element of $\mathcal{P}(q', \delta)$ for $q' \in M_1$ contained in the ball has norm between $\mathrm{Pac}_{\mathcal{P}}(L_\ell) + \delta$ and $\mathrm{Pac}_{\mathcal{P}}(L_\ell) - \delta$. Hence, $|\mathrm{Pac}_{\mathcal{P}}(L_n) - \mathrm{Pac}_{\mathcal{P}}(L_\ell)| \leq 2\delta \leq \epsilon$ for all $n \geq n(\delta)$.

□

Part II

Counting Cycles in the Datastreaming Model

Chapter 8

Counting Cycles in the Datastreaming Model

In this chapter, we develop an algebraic approach towards counting cycles of a fixed length in the data streaming model. The algorithm is based constructing complex-valued hash functions have posses certain desired cancellation properties.

Notation: Let $G = (V, E)$ be an undirected graph without self-loops and multiple edges. The set of vertices and edges are represented by $V[G]$ and $E[G]$ respectively. We will assume that $V[G] = \{1, \cdots, n\}$ and n is known in advance.

Given two directed graphs H_1 and H_2, we say that H_1 and H_2 are *homomorphic* if there is a mapping $i : V[H_1] \to V[H_2]$ such that $(u, v) \in E[H_1]$ if and only if $(i(u), i(v)) \in E[H_2]$. Furthermore, H_1 and H_2 are said to be *isomorphic* if the mapping i is a bijection.

For any graph H, we call a not necessarily induced subgraph H_1 of G an *occurrence* of H, if H_1 is isomorphic to H. We use $\#(H, G)$ to denote the number of occurrences of H in G. When G is the input graph, for simplicity we use $\#H$ to express $\#(H, G)$. Moreover, let C_ℓ be a cycle on ℓ edges.

8.1 A Review of Jowhari and Ghodsi's Algorithm

Our approach is best seen as a generalisation of the approach of Jowari and Ghodsi ([49]) to counting triangles in the data streaming model. Hence we start with a brief account of Jowhari and Ghodsi's algorithm in order to prepare the reader for our extension of their approach. Jowhari and Ghodsi estimate the number of triangles in a graph G. Let X be a $\{-1, +1\}$-valued random variable with expectation zero. They associate with every vertex w of G an instance $X(w)$ of X; the $X(w)$'s are 6-wise independent. They compute $Z = \sum_{\{u,v\} \in E[G]} X(u)X(v)$ and output $Z^3/6$ as the estimator for $\#C_3$.

Lemma 8.1.1 ([49]). $\mathbb{E}[Z^3] = 6 \cdot \#C_3$.

Proof. For any triple $T \in E^3[G]$ of edges and any vertex w of G, let $\deg_T(w)$ be the number of edges in T incident to w, then $\deg_T(w)$ is an integer no larger than 3. Also

$$\begin{aligned}\mathbb{E}[Z^3] &= \mathbb{E}\left[\left(\sum_{\{u,v\}\in E[G]} X(u)X(v)\right)^3\right] \\ &= \mathbb{E}\left[\sum_{T=(\{u_1,v_1\},\{u_2,v_2\},\{u_3,v_3\})\in E^3} X(u_1)X(v_1)X(u_2)X(v_2)X(u_3)X(v_3)\right]\end{aligned}$$

Let V_T be the set of vertices that are incident to the edges in T. Then

$$\mathbb{E}[Z^3] = \mathbb{E}\left[\sum_{T\in E^3}\prod_{w\in V_T} X(w)^{\deg_T(w)}\right]$$

By the 6-wise independence of the $X(w), w \in V$, we have

$$\mathbb{E}[Z^3] = \sum_{T\in E^3}\prod_{w\in V_T}\mathbb{E}\left[X(w)^{\deg_T(w)}\right] = \sum_{T\in E^3}\prod_{w\in V_T}\mathbb{E}\left[X^{\deg_T(w)}\right]$$

Since $\mathbb{E}\left[X^{\deg_T(w)}\right] = 1$ if $\deg_T(w)$ is even and $\mathbb{E}\left[X^{\deg_T(w)}\right] = 0$ if $\deg_T(w)$ is odd, we know that $\prod_{w\in V_T}\mathbb{E}\left[X^{\deg_T(w)}\right] = 1$ if and only if the edges in T form a triangle. Since each triangle is counted six times, we have $\mathbb{E}[Z^3] = 6\cdot \#C_3$. □ □

The crucial ingredients of the proof are (1) 6-wise independence guarantees that the expectation-operator can be pulled inside, and (2) random variable X is defined such that only vertices with even degree in T have nonzero expectation.

8.2 Algorithm Framework

We now generalize the algorithm in Section 8.1 and present an algorithm framework for counting general d-regular graphs. Suppose that H is a d-regular graph with k edges and we want to count the number of occurrences of H in G. The vertices of H are denoted by a, b and c etc, and the vertices of G are denoted by u, v and w, etc., respectively. We will equip the edges of H with an arbitrary orientation, as this is necessary for the further analysis. Therefore, each edge in H together with its orientation can be expressed as $\overrightarrow{ab}$ for some $a, b \in V[H]$. For simplicity and with slight abuse of notation we will use H to express such an oriented graph.

For each oriented edge $\overrightarrow{ab}$ in H our algorithm maintains a complex-valued variable $Z_{\overrightarrow{ab}}(G)$, which is initialized to zero. The variables are defined in terms of random variables $Y(w)$ and $X_c(w)$, where c is a node of H and w is a node of G. The random variables $Y(w)$ are instances of a random variable Y and the random variables $X_c(w)$ are instances of a random variable X. The range of both random variables is a finite subset of complex numbers. We will realize the random variables by hash functions from $V[G]$ to $\mathbb{C}$; this explains why we indicate the dependence on w by functional brackets.

We assume that the variables $X_c(w)$ and $Y(w)$ have sufficient independence as detailed below.

Our algorithm performs two basic steps: First, when an edge $e = \{u, v\} \in E[G]$ arrives, we update each variable $Z_{\overrightarrow{ab}}$ according to

$$Z_{\overrightarrow{ab}}(G) \leftarrow Z_{\overrightarrow{ab}}(G) + \big(X_a(u) \cdot X_b(v) + X_b(u) \cdot X_a(v)\big) \cdot Y(u) \cdot Y(v). \tag{8.1}$$

Second, when the number of occurrences of a graph H is required, the algorithm returns the real part of $Z/(\alpha \cdot \mathrm{aut}(H))$, where Z is defined via

$$Z := Z_H(G) = \prod_{\overrightarrow{ab} \in E[H]} Z_{\overrightarrow{ab}}(G), \tag{8.2}$$

α and $\mathrm{aut}(H)$ are constant numbers for any given H and will be determined later.

Remark 8.2.1. For simplicity, the algorithm above is only for the edge-insertion case. An edge deletion amounts to replacing '+' by '−' in (8.1).

Remark 8.2.2. The first step may be carried out in a distributed fashion, i. e., we have several processors each processing a subset of edges. In the second step the counts of the different processors are combined.

Theorem 8.2.3. *Let H be a d-regular graph with k edges. Let us assume that the random variables defined above satisfy the following two properties:*

1. *The random variables $X_c(w)$ and $Y(w)$, where $c \in V[H]$ and $w \in V[G]$, are instances of random variables X and Y, respectively. The random variables are $4k$-wise independent.*

2. *Let Z be any one of $X_c, c \in V[H]$ or Y. Then for any $1 \leq i \leq 2k$, $\mathbb{E}[Z^i] \neq 0$ if and only if $i = d$.*

Then $\mathbb{E}[Z_H(G)] = \alpha \cdot \mathrm{aut}(H) \cdot \#(H, G)$, where $\alpha = \left(\mathbb{E}\left[X^d\right] \mathbb{E}\left[Y^d\right]\right)^{2k/d} \in \mathbb{C}$ and $\mathrm{aut}(H)$ is the number of permutations and orientations of the edges in H such that the resulting graph is isomorphic to H.

The theorem above shows that $Z_H(G)$ is an unbiased estimator for any d-regular graph H, assuming that there exist random variables $X_c(w)$ and $Y(w)$ with certain properties. We will prove Theorem 8.2.3 at first, and then construct such random variables.

Proof of Theorem 8.2.3. We first introduce some notations. For a k-tuple $T = (e_1, \ldots, e_k) \in E^k[G]$, let $G_T = (V_T, E_T)$ be the induced multi-graph, i.e., G_T has edge multi-set $E_T = \{e_1, \ldots, e_k\}$. By definition, we have

$$\begin{aligned} Z_H(G) &= \prod_{\overrightarrow{ab} \in E[H]} Z_{\overrightarrow{ab}}(G) \\ &= \prod_{\overrightarrow{ab} \in E[H]} \left(\sum_{\{u,v\} \in E[G]} (X_a(u) \cdot X_b(v) + X_a(v) \cdot X_b(u)) \cdot Y(u) \cdot Y(v) \right). \end{aligned}$$

Since H has k edges, $Z_H(G)$ is a product of k terms and each term is a sum over all edges of G each with two possible orientations. Thus, in the expansion of $Z_H(G)$, any k-tuple $(e_1, \cdots, e_k) \in E^k[G]$ contributes 2^k different terms to $Z_H(G)$ and each term corresponds to a certain orientation of $(e_1, \cdots, e_k)$. Let $\overrightarrow{T} = (\overrightarrow{e_1}, \cdots, \overrightarrow{e_k})$ be an arbitrary orientation of $(e_1, \cdots, e_k)$, where $\overrightarrow{e_i} = \overrightarrow{u_i v_i}$. So the term in $Z_H(G)$ corresponding to $(\overrightarrow{e_1}, \cdots, \overrightarrow{e_k})$ is

$$\prod_{i=1}^{k} X_{a_i}(u_i) \cdot X_{b_i}(v_i) \cdot Y(u_i) \cdot Y(v_i) \ , \tag{8.3}$$

where (a_i, b_i) is the i-th edge of H and $\overrightarrow{u_i v_i}$ is the i-th edge in $\overrightarrow{T}$. We show that (8.3) is non-zero if and only if the graph induced by $\overrightarrow{T}$ is isomorphic to H (i. e. it also preserves the orientations of the edges).

For a vertex w of G and a vertex c of H, let

$$\theta_{\overrightarrow{T}}(c, w) = \left|\left\{i \mid (u_i = w \text{ and } a_i = c) \text{ or } (v_i = w \text{ and } b_i = c)\right\}\right| \ . \tag{8.4}$$

Thus for any $c \in V[H]$, $\sum_{w \in V_T} \theta_{\overrightarrow{T}}(c, w) = d$ since every vertex c of H appears in exactly d edges (a_i, b_i); recall that H is d-regular. Using the definition of $\theta_{\overrightarrow{T}}$, we rewrite (8.3) as

$$\left(\prod_{c \in V[H]} \prod_{w \in V_{\overrightarrow{T}}} X_c^{\theta_{\overrightarrow{T}}(c,w)}(w)\right) \cdot \left(\prod_{w \in V_{\overrightarrow{T}}} Y^{\deg_{\overrightarrow{T}}(w)}(w)\right),$$

where $\deg_{\overrightarrow{T}}(w)$ is the number of edges in $\overrightarrow{T}$ incident to w. Therefore

$$\begin{aligned} & Z_H(G) \\ = & \sum_{\substack{e_1, \cdots, e_k \\ e_i \in E[G]}} \sum_{\overrightarrow{T} = (\overrightarrow{e_1}, \cdots, \overrightarrow{e_k})} \left(\prod_{c \in V[H]} \prod_{w \in V_{\overrightarrow{T}}} X_c^{\theta_{\overrightarrow{T}}(c,w)}(w)\right) \cdot \left(\prod_{w \in V_{\overrightarrow{T}}} Y^{\deg_{\overrightarrow{T}}(w)}(w)\right), \end{aligned}$$

where the first summation is over all the k-tuples of edges in $E[G]$ and the second summation is over all their possible orientations. Since each term of Z_H is the product of $4k$ random variables, which by assumption are $4k$-wise independent, we infer by linearity of expectation that

$$\begin{aligned} & \mathbb{E}[Z_H(G)] \\ = & \mathbb{E}\left[\sum_{\substack{e_1, \cdots, e_k \\ e_i \in E[G]}} \sum_{\overrightarrow{T} = (\overrightarrow{e_1}, \cdots, \overrightarrow{e_k})} \left(\prod_{c \in V[H]} \prod_{w \in V_{\overrightarrow{T}}} X_c^{\theta_{\overrightarrow{T}}(c,w)}(w)\right) \cdot \left(\prod_{w \in V_{\overrightarrow{T}}} Y^{\deg_{\overrightarrow{T}}(w)}(w)\right)\right] \\ = & \sum_{\substack{e_1, \cdots, e_k \\ e_i \in E[G]}} \sum_{\overrightarrow{T} = (\overrightarrow{e_1}, \cdots, \overrightarrow{e_k})} \prod_{c \in V[H]} \prod_{w \in V_{\overrightarrow{T}}} \mathbb{E}\left[X^{\theta_{\overrightarrow{T}}(c,w)}\right] \cdot \prod_{w \in V_{\overrightarrow{T}}} \mathbb{E}\left[Y^{\deg_{\overrightarrow{T}}(w)}\right] \ . \end{aligned}$$

Let

$$\alpha(\overrightarrow{T}) := \prod_{c \in V[H]} \prod_{w \in V_{\overrightarrow{T}}} \mathbb{E}\left[X^{\theta_{\overrightarrow{T}}(c,w)}\right] \cdot \prod_{w \in V_{\overrightarrow{T}}} \mathbb{E}\left[Y^{\deg_{\overrightarrow{T}}(w)}\right] \ .$$

We will next show that $\alpha(\overrightarrow{T})$ is either zero or a nonzero constant independent of $\overrightarrow{T}$. The latter is the case if and only if G_T is an occurrence of H in G.

We have $\mathbb{E}\left[X^i\right] \neq 0$ if and only if $i = d$ or $i = 0$. Therefore for any $\overrightarrow{T}$ and $c \in V[H]$, $\prod_{w \in V_{\overrightarrow{T}}} \mathbb{E}[X^{\theta_{\overrightarrow{T}}(c,w)}] \neq 0$ if and only if $\theta_{\overrightarrow{T}}(c,w) \in \{0, d\}$ for all w. Since $\sum_w \theta_{\overrightarrow{T}}(c,w) = \deg_H(c) = d$, there must be a unique vertex $w \in V_{\overrightarrow{T}}$ such that $\theta_{\overrightarrow{T}}(c,w) = d$. Define $\varphi : V[H] \to V_{\overrightarrow{T}}$ as $\varphi(c) = w$. Then φ is a homomorphism and

$$\prod_{c \in V[H]} \prod_{w \in V_{\overrightarrow{T}}} \mathbb{E}\left[X^{\theta_{\overrightarrow{T}}(c,w)}\right] = \prod_{c \in V[H]} \mathbb{E}\left[X^d\right] = \mathbb{E}\left[X^d\right]^{|V[H]|}.$$

Since $\mathbb{E}[Y^i] \neq 0$ if and only if $i = d$ or $i = 0$, so for any $\overrightarrow{T}$, $\prod_{w \in V_{\overrightarrow{T}}} \mathbb{E}[Y^{\deg_{\overrightarrow{T}}(w)}] \neq 0$ if and only if every vertex $w \in V_{\overrightarrow{T}}$ has degree d in the graph with edge set T. Thus $|V_{\overrightarrow{T}}| = 2k/d = |V[H]|$, which implies that φ is an isomorphism mapping.

We have now shown that $\alpha(\overrightarrow{T})$ is either zero or the nonzero constant

$$\alpha = \left(\mathbb{E}\left[X^d\right] \mathbb{E}\left[Y^d\right]\right)^{2k/d}.$$

The latter is the case if and only if $G_{\overrightarrow{T}}$ is an occurrence of H in G. Let $(G_{\overrightarrow{T}} \equiv H)$ be the indicator expression that is one if $G_{\overrightarrow{T}}$ and H are isomorphic and zero otherwise. Then

$$\mathbb{E}[Z_H(G)] = \sum_{\substack{e_1, \cdots, e_k \\ e_i \in E[G]}} \sum_{\overrightarrow{T} = (\overrightarrow{e_1}, \cdots, \overrightarrow{e_k})} \alpha(\overrightarrow{T}) \cdot \left(G_{\overrightarrow{T}} \equiv H\right) = \alpha \cdot \mathrm{aut}(H) \cdot \#(H, G) \ .$$

□ □

For the case of cycles, we have $\mathrm{aut}(H) = 2k$. We turn to construct hash functions needed in Theorem 8.2.3. The basic idea is to choose a $8k$-wise independent hash function $h : D \to \mathbb{C}$ and map the values in D to complex numbers with certain properties. We first show a simple lemma about roots of polynomials of a simple form.

Lemma 8.2.4. *For positive interger r, let $P_r(z) = 2 + z^r$ and $z_j = 2^{1/j} \cdot \mathrm{e}^{\frac{\pi i}{j}}$. The complex number z_j is a root of the polynomial $P_r(z)$ if and only if $j = r$.*

Proof. We first verify that z_r is a root of the polynomial $P_r(z)$: since $z_r^r = 2 \cdot \mathrm{e}^{\pi \cdot i} = -2$, we have $z_r^r + 2 = 0$. To show the converse, we consider z_j^r for $r \neq j$ and verify that $|z_j^r| = \left|2^{r/j} \mathrm{e}^{\frac{\pi \cdot i \cdot r}{j}}\right| = 2^{r/j}$. Since $2^{r/j} \neq 2$ if $j \neq r$, the claim follows. □ □

Let z_j as in Lemma 8.2.4 and define random variable H_j as

$$H_j = \begin{cases} 1, & \text{with probability } 2/3, \\ z_j, & \text{with probability } 1/3. \end{cases} \tag{8.5}$$

Then $\mathbb{E}[H_j^\ell] = \left(2 + z_j^\ell\right)/3 = P_\ell(z_j)/3$ which is nonzero if $j \neq \ell$.

Theorem 8.2.5. *For positive integers d and k, let*

$$H = \prod_{1 \le j \le 2k, j \ne d} H_j$$

where the H_j are independent. For all integers ℓ between 1 and $2k$, $\mathbb{E}[H^\ell] \neq 0$ if and only if $d = \ell$.

Proof. By independence, $\mathbb{E}[H^\ell] = \prod_{1 \le j \le 2k, j \ne d} \mathbb{E}[H_j^\ell]$. This product is nonzero if ℓ is different from all j that are distinct from d, i. e., $\ell = d$. □ □

8.3 Proof of the Main Theorem

Now we bound the space of the algorithm for the case of cycles of arbitrary length. The basic idea is to use the second moment method on the complex-valued random variable Z. We first note a couple of lemmas that turn out to be useful: the first lemma is a generalization of Chebyshev's inequality for a complex-valued random variable and the second lemma is an upper bound on the number of closed walks of a given length in terms of the number of edges of the graph. Recall that the conjugate of a complex number $z = a + ib$ is denoted by $\overline{z} := a - ib$.

Lemma 8.3.1. *Let X be a complex-valued random variable with finite support and let $t > 0$. We have that*

$$\Pr[|X - \mathbb{E}[X]| \ge t \cdot |\mathbb{E}[X]|] \le \frac{\mathbb{E}[X\overline{X}] - \mathbb{E}[X]\overline{\mathbb{E}[X]}}{t^2 |\mathbb{E}[X]|^2} \quad .$$

Proof. Since $|X - \mathbb{E}[X]|^2 = (X - \mathbb{E}[X])(\overline{X - \mathbb{E}[X]})$ is a positive-valued random variable, we apply Markov's inequality to obtain

$$\begin{aligned}\Pr[|X - \mathbb{E}[X]| \ge t \cdot |\mathbb{E}[X]|] &= \Pr\left[|X - \mathbb{E}[X]|^2 \ge t^2 \cdot |\mathbb{E}[X]|^2\right] \\ &\le \frac{\mathbb{E}[(X - \mathbb{E}[X])\overline{(X - \mathbb{E}[X])}]}{t^2 |\mathbb{E}[X]|^2} \quad .\end{aligned}$$

Expanding $\mathbb{E}[(X - \mathbb{E}[X])\overline{(X - \mathbb{E}[X])}]$ we obtain that

$$\begin{aligned}\mathbb{E}[(X - \mathbb{E}[X])\overline{(X - \mathbb{E}[X])}] &= \mathbb{E}[X\overline{X}] - \mathbb{E}[X\overline{\mathbb{E}[X]}] - \mathbb{E}[\overline{X}\mathbb{E}[X]] + \overline{\mathbb{E}[X]}\mathbb{E}[X] \\ &= \mathbb{E}\left[X\overline{X}\right] - [X]\overline{\mathbb{E}[X]} \quad .\end{aligned}$$

The last equality uses the linearity of expectation and that $\mathbb{E}[\overline{X}] = \overline{\mathbb{E}[X]}$. □ □

We now show an upper bound on the number of closed walks of a given length in a graph. This upper bound will control the space requirement of the algorithm.

Lemma 8.3.2. *Let G be an undirected graph with n vertices and m edges. Then the number of closed walks W_k with length k in G is at most $\frac{2^{k/2-1}}{k} \cdot m^{k/2}$.*

Proof. Let A be the adjacency matrix of G with eigenvalues $\lambda_1, \cdots, \lambda_n$. Since G is undirected, A is real symmetric and each eigenvalue λ_i is a real number. Then $W_k = \frac{1}{2k} \cdot \sum_{i=1}^{n} (A^k)_{ii}$ where for a matrix M, M_{ij} is the ij-th entry of the matrix. Because $\sum_{i=1}^{n} (A^k)_{ii} = \mathrm{tr}(A^k) = \sum_{i=1}^{n} \lambda_i^k \le \sum_{i=1}^{n} |\lambda_i|^k$ and $\left(\sum_{i=1}^{n} |\lambda_i|^k\right)^{1/k} \le (\sum_{i=1}^{n} |\lambda_i|^2)^{1/2} = (2m)^{1/2}$ for any $k \ge 2$, we have $W_k \le \frac{1}{2k} \cdot (\sum_{i=1}^{n} |\lambda_i|^2)^{k/2} = \frac{2^{k/2-1}}{k} \cdot m^{k/2}$. □ □

Corollary 8.3.3. *Let G be a graph on m edges and $\mathcal{H}$ be a set of subgraphs of G such that every $H \in \mathcal{H}$ has properties: (1) H has k edges, where k is a constant. (2) Each connected-component of H is an Eulerian circuit. Then $|\mathcal{H}| = O(m^{k/2})$.*

Proof. Fix an integer $r \in \{1, \ldots, k\}$ and consider graphs in $\mathcal{H}$ that have r connected components. By Lemma 8.3.2, the number of such graphs is at most

$$\sum_{\substack{k_1,\cdots,k_r \\ k_1+\cdots+k_r=k}} \prod_{i=1}^{r} W_{k_i} \le \sum_{\substack{k_1,\cdots,k_r \\ k_1+\cdots+k_r=k}} \prod_{i=1}^{r} \frac{2^{k_i/2-1} \cdot m^{k_i/2}}{k_i} \le f(k) \cdot (2m)^{k/2},$$

where $f(k)$ is a function of k. Because there are at most k choices of r, we have $|\mathcal{H}| = O(m^{k/2})$. □ □

Observe that the expansion of $\mathbb{E}[Z_H(G)\overline{Z_H(G)}]$ consists of m^{2k} terms and the modulus of each term is upper bounded by a constant. So a naïve upper bound for $\mathbb{E}[Z_H(G)\overline{Z_H(G)}]$ is $O(m^{2k})$. Now we only focus on the case of cycles and use the "cancellation" properties of the random variables to get a better bound for $\mathbb{E}[Z_H(G)\overline{Z_H(G)}]$.

Theorem 8.3.4. *Let H be a cycle C_k with an arbitrary orientation and suppose that the following properties are satisfied:*

1. *The random variables $X_c(w)$ and $Y(w)$, where $c \in V[H]$ and $w \in V[G]$ are $8k$-wise independent.*
2. *Let Z be any one of $X_c, c \in V[H]$ or Y. Then for any $1 \le i \le 2k$, $\mathbb{E}[Z^i] \neq 0$ if and only if $i = 2$.*

Then $\mathbb{E}[Z_H(G)\overline{Z_H(G)}] = O(m^k)$.

Proof. By the definition of $Z_H(G)$ we express $Z_H(G)\overline{Z_H(G)}$ as

$$\sum_{\substack{\overrightarrow{T_1}=(\overrightarrow{e_1},\cdots,\overrightarrow{e_k}) \\ \overrightarrow{T_2}=(\overrightarrow{e'_1},\cdots,\overrightarrow{e'_k}) \\ e_i,e'_i\in E[G]}} \left(\prod_{\substack{c\in V[H] \\ w\in V_{\overrightarrow{T_1}}}} X_c(w)^{\theta_{\overrightarrow{T_1}}(c,w)} \right) \cdot \left(\prod_{w\in V_{\overrightarrow{T_1}}} Y(w)^{\deg_{\overrightarrow{T_1}}(w)} \right) \cdot$$

$$\left(\prod_{\substack{c\in V[H] \\ w\in V_{\overrightarrow{T_2}}}} \overline{X_c(w)}^{\theta_{\overrightarrow{T_2}}(c,w)} \right) \cdot \left(\prod_{w\in V_{\overrightarrow{T_2}}} \overline{Y(w)}^{\deg_{\overrightarrow{T_2}}(w)} \right),$$

where the function $\theta_{\overrightarrow{T}}(\cdot,\cdot)$ is defined in (8.4). Using the linearity of expectations and the $8k$-wise independence of the random variables $X_c(w)$ and $Y(w)$, we obtain

$$\mathbb{E}\left[Z_H(G)\overline{Z_H(G)}\right] = \sum_{\substack{\overrightarrow{T_1}=(\overrightarrow{e_1},\cdots,\overrightarrow{e_k})\\ \overrightarrow{T_2}=(\overrightarrow{e'_1},\cdots,\overrightarrow{e'_k})\\ e_i,e'_i\in E[G]}} Q_{\overrightarrow{T_1},\overrightarrow{T_2}},$$

where

$$Q_{\overrightarrow{T_1},\overrightarrow{T_2}} = \left(\prod_{c\in V[H]}\prod_{w\in V_{\overrightarrow{T_1}}\cup V_{\overrightarrow{T_2}}} \mathbb{E}\left[X_c(w)^{\theta_{\overrightarrow{T_1}}(c,w)}\overline{X_c(w)}^{\theta_{\overrightarrow{T_2}}(c,w)}\right]\right)\cdot$$
$$\left(\prod_{w\in V_{\overrightarrow{T_1}}\cup V_{\overrightarrow{T_2}}} \mathbb{E}\left[Y(w)^{\deg_{\overrightarrow{T_1}}(w)}\overline{Y(w)}^{\deg_{\overrightarrow{T_2}}(w)}\right]\right).$$

For any $c\in V[H]$ and $w\in V_{\overrightarrow{T_1}}\cup V_{\overrightarrow{T_2}}$, we write

$$R_{\overrightarrow{T_1},\overrightarrow{T_2}}(c,w) = \mathbb{E}\left[X_c(w)^{\theta_{\overrightarrow{T_1}}(c,w)}\overline{X_c(w)}^{\theta_{\overrightarrow{T_2}}(c,w)}\right].$$

Let $R_{\overrightarrow{T_1},\overrightarrow{T_2}} = \prod_{c\in V[H]}\prod_{w\in V_{\overrightarrow{T_1}}\cup V_{\overrightarrow{T_2}}} R_{\overrightarrow{T_1},\overrightarrow{T_2}}(c,w)$. Then

$$Q_{\overrightarrow{T_1},\overrightarrow{T_2}} = R_{\overrightarrow{T_1},\overrightarrow{T_2}}\cdot\prod_{w\in V_{\overrightarrow{T_1}}\cup V_{\overrightarrow{T_2}}} \mathbb{E}\left[Y(w)^{\deg_{\overrightarrow{T_1}}(w)}\overline{Y(w)}^{\deg_{\overrightarrow{T_2}}(w)}\right].$$

We claim that if the term $Q_{\overrightarrow{T_1},\overrightarrow{T_2}}\neq 0$, then every vertex in $V_{\overrightarrow{T_1}}\cup V_{\overrightarrow{T_2}}$ has even degree in the undirected sense. First, we show that using this claim we can finish the proof of the theorem. Note that $\mathbb{E}[Z_H(G)\overline{Z_H(G)}] = \sum_{G_{\overrightarrow{T_1},\overrightarrow{T_2}}\in\mathcal{E}_{2k}} Q_{\overrightarrow{T_1},\overrightarrow{T_2}}$ where $\mathcal{E}_{2k}$ is the set of directed subgraphs of G on $2k$ edges with every vertex having even degree in the undirected sense. Observing that the undirected graph defined by $G_{\overrightarrow{T_1},\overrightarrow{T_2}}$ is a Eulerian circuit, by Corollary 8.3.3 we get $\mathbb{E}[Z_H(G)\overline{Z_H(G)}]\leq\sum_{G_{\overrightarrow{T_1},\overrightarrow{T_2}}\in\mathcal{E}_{2k}}|Q_{\overrightarrow{T_1},\overrightarrow{T_2}}|\leq c\cdot m^k$. Note that an upper bound for the constant c is $\max_{G_{\overrightarrow{T_1},\overrightarrow{T_2}}\in\mathcal{E}_{2k}}|Q_{\overrightarrow{T_1},\overrightarrow{T_2}}|$.

Let us now prove that $Q_{\overrightarrow{T_1},\overrightarrow{T_2}}\neq 0$ implies that every vertex in $V_{\overrightarrow{T_1}}\cup V_{\overrightarrow{T_2}}$ has even degree in the undirected sense. We first make the following observations: For any vertex c of C_k and w in $V_{\overrightarrow{T_1}}\cup V_{\overrightarrow{T_2}}$ we have: $\mathbb{E}\left[X_c^i(w)\right]\neq 0$ if and only if $i=2$. After expanding $Z_H(G)$ and $\overline{Z_H(G)}$, $X_c(\cdot), c\in V[H]$ appears twice in each term, so we have $\sum_{w\in V_{\overrightarrow{T_1}}\cup V_{\overrightarrow{T_2}}}\theta_{\overrightarrow{T_1}}(c,w)+\theta_{\overrightarrow{T_2}}(c,w)=4$. Consider a subgraph $G_{\overrightarrow{T_1},\overrightarrow{T_2}}$ on $2k$ edges such that $R_{\overrightarrow{T_1},\overrightarrow{T_2}}\neq 0$. Assume for the sake of contradiction that $G_{\overrightarrow{T_1},\overrightarrow{T_2}}$ has a vertex w of odd degree. This implies that there is a vertex $c\in C_k$ such that $\theta_{\overrightarrow{T_1}}(c,w)+\theta_{\overrightarrow{T_2}}(c,w)$ is either one or three. However $\theta_{\overrightarrow{T_1}}(c,w)+\theta_{\overrightarrow{T_2}}(c,w)$ cannot be one since in this case both $R_{\overrightarrow{T_1},\overrightarrow{T_2}}$ and $Q_{\overrightarrow{T_1},\overrightarrow{T_2}}$ must vanish. Now consider the case where $\theta_{\overrightarrow{T_1}}(c,w)+\theta_{\overrightarrow{T_2}}(c,w)=3$. This means that $R_{\overrightarrow{T_1},\overrightarrow{T_2}}(c,w)$ is either $\mathbb{E}[X_c^2(w)\overline{X_c(w)}]$ or the symmetric variant $\mathbb{E}\big[X_c(w)\overline{X_c(w)}^2\big]$.

Assume that $R_{\overrightarrow{T_1},\overrightarrow{T_2}}(c,w) = \mathbb{E}[X_c^2(w)\overline{X_c(w)}]$. Since $\sum_{w\in V_{\overrightarrow{T_1}}\cup V_{\overrightarrow{T_2}}} \theta_{\overrightarrow{T_1}}(c,w)+\theta_{\overrightarrow{T_2}}(c,w) = 4$, there must be a vertex $w' \neq w$ in $V_{\overrightarrow{T_1}} \cup V_{\overrightarrow{T_2}}$ such that $R_{\overrightarrow{T_1},\overrightarrow{T_2}}(c,w') = \mathbb{E}[\overline{X_c(w')}]$. This implies that $R_{\overrightarrow{T_1},\overrightarrow{T_2}}$ vanishes and hence $Q_{\overrightarrow{T_1},\overrightarrow{T_2}}$ must also vanish, which leads to a contradiction. □

□

We are now ready to prove the main result of this chapter.

Theorem 8.3.5. *Let G be a graph with n vertices and m edges. For any k, there is an algorithm using S bits of space to (ε,δ)-approximate the number of occurrences of C_k in G provided that $S = \Omega\big(\frac{1}{\varepsilon^2}\cdot\frac{m^k}{(\#C_k)^2}\cdot\log n\cdot\log\frac{1}{\delta}\big)$. The algorithm works in the turnstile model.*

Proof. First, observe that

$$\frac{\mathbb{E}[Z_H(G)\overline{Z_H(G)}]-\mathbb{E}^2[Z_H(G)]}{|\mathbb{E}[Z_H(G)]|^2} \leq \frac{\mathbb{E}[Z_H(G)\overline{Z_H(G)}]}{|\mathbb{E}[Z_H(G)]|^2}.$$

We run s parallel and independent copies of our estimator, and take the average value $Z^* = \frac{1}{s}\sum_{i=1}^{s} Z_i$, where each Z_i is the output of the i-th instance of the estimator. Therefore $\mathbb{E}[Z^*] = \mathbb{E}[Z_H(G)]$ and

$$\mathbb{E}[Z^*\overline{Z}^*] - |\mathbb{E}\left[Z^*\right]|^2 = \frac{1}{s}\left(\mathbb{E}[Z_H(G)\overline{Z_H(G)}] - |\mathbb{E}[Z_H(G)]|^2\right).$$

By Chebyshev's inequality (Lemma 8.3.1), we have

$$\Pr\left[|Z^*-\mathbb{E}[Z^*]| \geq \varepsilon\cdot|\mathbb{E}[Z^*]|\right] \leq \frac{\mathbb{E}[Z_H(G)\overline{Z_H(G)}]-\mathbb{E}[Z_H(G)]\overline{\mathbb{E}[Z_H(G)]}}{s\cdot\varepsilon^2\cdot|\mathbb{E}[Z_H(G)]|^2}.$$

Observe that

$$\mathbb{E}[Z_H(G)\overline{Z_H(G)}]-\mathbb{E}[Z_H(G)]\overline{\mathbb{E}[Z_H(G)]} \leq \mathbb{E}[Z_H(G)\overline{Z_H(G)}] = O(m^k).$$

By choosing $s = O\left(\frac{1}{\varepsilon^2}\cdot\frac{m^k}{(\#C_k)^2}\right)$, we get $\Pr\left[|Z^*-\mathbb{E}\left[Z^*\right]| \geq \varepsilon\cdot|\mathbb{E}[Z^*]|\right] \leq 1/3$.

The probability of success can be amplified to $1-\delta$ by running in parallel $O\left(\log\frac{1}{\delta}\right)$ copies of the algorithm and outputting the median of those values.

Since storing each random variable requires $O(\log n)$ space and the number of random variables used in each trial is $O(1)$, so the overall space complexity is as claimed. □ □

8.4 Future Directions

We now briefly discuss some natural extensions of our results. We can count other graphs by constructing hash functions that take values from non-commutative algebras such as Clifford algebras. Interestingly, this approach of Clifford algebras has been successfully used in approximating the permanent [27]. Another direction is to obtain lower bounds for counting various subgraphs in the data streaming model.

Bibliography

[1] T. Akutsu, *On determining the congruence of point sets in d dimensions*, Computational Geometry, 247–256, 1998.

[2] M. Aigner, G. Ziegler, and K. Hofmann. *Proofs from THE BOOK*. Springer, 2010.

[3] O. Amini and M. Manjunath, *Riemann-Roch for Sub-lattices of the Root Lattice A_n*, Electronic Journal of Combinatorics (**17(1)**), 2010.

[4] C. Ambühl, S. Chakraborty, and Bernd Gärtner. *Computing largest common point sets under approximate congruence*, in the proceedings of the European Symposium on Algorithms, 52–63, 2000.

[5] I. V. Artamkin, *Discrete Torelli Theorem*, Matematicheski, Sbornik (**197:8**), 3–16, 2006.

[6] V. Arvind, R. Beigel, and A. Lozano, *The complexity of modular graph automorphism*, SIAM Journal on Computing, 30 **(4)**, 1299–1320, 2000.

[7] A. Asadi and S. Backman, *Chip-Firing and Riemann-Roch Theory for Directed Graphs*, in arxiv at arXiv:1012.0287v1, 2010.

[8] L. Babai, *Character sums, Weil's Estimates and Paradoxical Tournaments.* http://people.cs.uchicago.edu/ laci/reu02.dir/paley.pdf.

[9] L. Babai, D. Y. Grigoryev, and D. M. Mount. *Isomorphism of graphs with bounded eigenvalue multiplicity*. In Proceedings of the fourteenth annual ACM symposium on Theory of computing, 310–324, 1982.

[10] R. Bacher, P. Harpe and T. Nagnibeda, *The lattice of integral flows and the lattice of integral cuts on a finite graph*, Bulletin de la Société Mathématiques de France, 125 **(2)**, 167–198, 1997.

[11] M. Baker, *Specialization of Linear Systems from Curves to Graphs*, Algebra and Number Theory 2**(6)** (2008), 613–653.

[12] M. Baker and S. Norine, *Riemann-Roch and Abel-Jacobi Theory on a Finite Graph*, Advances in Mathematics 215 **(2)**, 766–788, 2007.

[13] M. Baker and F. Shokrieh, *Chip-Firing Games and Potential Theory on Graphs, and Spanning trees*, in arxiv at arXiv:1107.1313, 2011.

[14] S. Basu, R. Pollack, and M.-F. Roy *Algorithms in Real Algebraic Geometry*, Springer, 2nd edition, 2006.

[15] A. Barvinok, *Computing the Ehrhart quasi-polynomial of a rational simplex*, Mathematics of Computation, 75, 1449–1466, 2006.

[16] N. Biggs, *Algebraic graph theory*, Cambridge University Press, 1993.

[17] B. Benson, D. Chakrabarty and P. Tetali: *G-parking functions, acyclic orientations and spanning trees*, Discrete Mathematics (2010).

[18] A. Björner, L. Lovász and P.W. Shor, *Chip-firing games on graphs*, European Journal of Combinatorics 12 (**4**), 283–291, 1991.

[19] P. Brass and C. Knauer. *Testing the congruence of d-dimensional point sets*. In Proceedings of the sixteenth annual symposium on Computational geometry, 310–314, 2000.

[20] I. Bordino and D. Donato and A. Gionis and S. Leonardi, *Mining Large Networks with Subgraph Counting*, Proceedings of the 8th IEEE International Conference on Data Mining, 737–742.

[21] N. Bourbaki, *Lie Groups and Lie Algebras*, Springer, 1989.

[22] L. S. Buriol, G. Frahling, S. Leonardi, A. Marchetti-Spaccamela and C. Sohler *Counting triangles in data streams*, Proceedings of the 25th ACM Symposium on Principles of Database Systems, 2006, 253–262.

[23] L. Caporaso, *Algebraic and combinatorial Brill-Noether theory*, in arxiv at arXiv:1106.1140v1, 2011.

[24] L. Caporaso and F. Viviani, *Torelli theorem for graphs and tropical curves*, Duke Mathematical Journal, 153 (**(1)**, (2010), 129-171.

[25] L.P. Chew and R.L.S Dyrsdale III, *Voronoi Diagrams based on Convex Distance Functions*, Proceedings of the First Annual Symposium on Computational Geometry, 235–244, 1985.

[26] D. Chebikin and P. Pylyavskyy, *A family of bijections between G-parking functions and spanning trees*, Journal of Combinatorial Theory Series A, 110(**1**), 31–41, 2005.

[27] S. Chien, L. Rasmussen and A. Sinclair, *Clifford algebras and approximating the permanent*, Journal of Computer and System Sciences, 67**(2)**, 263–290, 2003.

[28] F.R.K Chung, *Spectral Graph Theory*, American Mathematical Society, 1997.

[29] F. Cools, J. Draisma, S. Payne and E. Robeva, *A tropical proof of the Brill-Noether Theorem*, in arxiv at arXiv:1001.2774v2, 2011.

[30] J.H Conway and N.J.A Sloane, *The cell structures of certain lattices*, Miscellanea Mathematica, Springer Verlag , 71–107, 1991.

[31] J.H Conway and N.J.A Sloane, *Sphere Packings, Lattices and Groups*, Springer Verlag, 1998.

[DG03] S. Dasgupta and A. Gupta. *An elementary proof of a theorem of Johnson and Lindenstrauss*. Random Struct. Algorithms, 22 **(1)**, 60–65, 2003.

[32] D. Dhar, *Self-organised critical state of sandpile automation models*, Physical Review Letters, 64**(14)**: 1613–1616, 1990.

[33] M. Develin and B. Sturmfels, *Tropical Convexity*, Documenta Mathematica 9, 1–27, 2004.

[34] A. M. Duval, C. J. Klivans, and J. L. Martin, *Simplicial Matrix-Tree Theorems*, in arxiv at arXiv:0802.2576v2 [math.CO], 2010.

[35] H. Edelsbrunner. *Geometry and Topology for Mesh Generation*, Cambridge Univ. Press, England, 89–104, 2001.

[36] M. Eichler, *Quadratische Formen und Orthogonal Gruppen*, Springer Verlag, 1952.

[37] J. Flum and M. Grohe, *The Parameterized Complexity of Counting Problems*, SIAM Journal on Computing, 33 **(4)**, 892–922, 2004.

[38] S. Ganguly, *Estimating Frequency Moments of Data Streams Using Random Linear Combinations*, in APPROX-RANDOM, 369–380, 2004.

[39] A. Gathmann and M. Kerber *A Riemann-Roch theorem in tropical geometry*, Mathematische Zeitschrift, 259, 217-230, 2008.

[40] J. Giesen and M. John, *The Flow Complex: A Data Structure For Geometric Modeling*, Computational Geometry: Theory and Applications, 39(**3**), 178–190, 2008.

[41] C. Godsil and G. Royle, *Algebraic Graph Theory*, Springer-Verlag, 2001.

[42] I. Gohberg, P. Lancaster, and L. Rodman, *Invariant Subspaces of Matrices with Applications*, SIAM Classics in Applied Mathematics. SIAM, 2006.

[43] D. Y. Grigor'ev, *Two reductions of graph isomorphism to problems on polynomials. Journal of Mathematical Sciences*, 20**(4)**, 2296–2298, 1982.

[44] P. M. Gruber, C. G. Lekkerkerker, *Geometry of numbers*, North-Holland Mathematical Library, Elsevier, Second Edition, 1987.

[45] J. van den Heuvel, *Algorithmic aspects of a chip-firing game*, Combinatorics, Probability and Computing 10 **(6)**, 505-529, 2001.

[46] J. Hladký, D. Kräl, S. Norine, *Rank of divisors on tropical curves*, in arxiv at arXiv:0709.4485v3, 2007.

[47] R. A. Horn and C. R. Johnson, *Matrix Analysis*, Cambridge University Press, 405, 1985.

[48] S. Hoory, N. Linial and A. Wigderson, *Expander graphs and their applications*, Bulletin of the AMS, 43, 439–561, 2006.

[49] H. Jowhari and M. Ghodsi, *New Streaming Algorithms for Counting Triangle Graphs*, Proceedings of 11th Annual International Computing and Combinatorics Conference, 2005.

[50] V. Kaibel and A. Schwartz. *On the complexity of polytope isomorphism problems.* Graphs and Combinatorics, 19**(2)**, 215–230, 2003.

[51] R. Kannan, *Minkowski's Convex Body Theorem and Integer Programming*, Mathematics of Operations Research, 12 **(3)**, 415–440, 1987.

[52] R. Kannan, *Algorithmic Geometry of Numbers*, Annual Review of Computer Science, 2, 231–267, 1987.

[53] N. Karmarkar and R. Karp and R. Lipton and L. Lovasz and M. Luby, *A Monte-Carlo Algorithm for Estimating the Permanent*, SICOMP: SIAM Journal on Computing **22**, 284-293, 1993.

[54] J. Köbler, U. Schöning, and J. Torán. *The Graph Isomorphism Problem: Its Structural Complexity*, Progress in Theoretical Computer Science. Birkhauser, 1993.

[55] S. Lang. *Algebra*, Springer, 2002.

[56] S. Leonardi and C. Sohler, *Estimating Clustering Indexes in Data Stream*, Proceedings of the 15th Annual European Symposium, Springer, Lecture Notes in Computer Science, 618–632, 2007.

[57] J. A. de Loera, J. Rambau and F. Santos, *Triangulations: Structures and Algorithms*, 2010.

[58] D. Lorenzini, *Two-Variable-Zeta-Functions On Graphs and Riemann-Roch theorems*, available at "http:www.math.uga.edu/ lorenz/ZetaFunctionGraphs.pdf", 2011.

[59] A. Lubiw. *Some NP-complete problems similar to Graph Isomorphism*, SIAM Journal of Computing, 10, 11–21, 1981.

[60] M. Manjunath, K. Mehlhorn, K. Panagiotou and H. Sun, *Approximate Counting of Cycles in Streams*, in the proceedings of European Symposion of Algorithms, 2011.

[61] M. Manjunath and V. Sharma, *Applications of Dimensionality Reduction and Exponential Sums to Graph Automorphism*, Theoretical Computer Science, 412 **(29)**, 2011.

[62] J. Martinet *Perfect Lattices in Euclidean Spaces*, Springer, 2003.

[63] G. Mikhailin and I. Zharkov, *Tropical Curves, their Jacobians and theta functions*, Curves and Abelian varieties, Contemporary Mathematics, 465, 203–230, 2008.

[64] B. Mohar and I. Rivin, *Simplices and spectra of graphs*, Discrete and Computational Geometry, 43 **(3)**, 516–521, 2010.

[65] K. Mehlhorn and P. Sanders. *Algorithms and Data Structures*, Springer, 2008.

[66] S. Muthukrishnan, *Data Streams: Algorithms and Applications*, Foundations and Trends in Theoretical Computer Science, 1 **(2)**, 2005.

[67] G. Mikhalkin and I. Zharkov, *Tropical curves, their Jacobians and Theta functions*, in arxiv at arXiv:math/0612267v2, 2007.

[68] R. Nellamani, S. Dash and R. G. Baraniuk, *On Nearly Orthogonal Lattice Bases and Random Lattices*, SIAM Journal on Discrete Mathematics, 21 **(1)**, 199–219, 2007.

[69] A. McGregor, *Open Problems in Data Streams and Related Topics*, IITK Workshop on Algoriths For Data Sreams, 2006.

[70] I. Peeva and B. Sturmfels, *Generic Lattice Ideals*, Journal of the American Mathematical Society, 11**(2)**, 363–373, 1998.

[71] A. Postnikov and B. Shapiro, *Trees, parking functions, syzygies, and deformations of monomial ideals*, Transactions of the American Mathematical Society, 356 **(8)**, 3109–3142, 2004.

[72] V. Y. Protasov. *Graph isomorphism and equality of simplices*. Mathematical Notes, 85(5):724–732, 2009. Original Russian Text published in Mathematicheskie Zametki, 85**(5)**, 758–767, 2009.

[73] H. Radjavi and P. Rosenthal, *Invariant Subspaces*, Springer-Verlag, 1973.

[74] H. Scarf, *The structure of the complex of maximal lattice point free bodies for a matrix of size* $n+1 \times n$, Bolyai Society Mathematical Studies, **19**, 1–32, 2008.

[75] R. Schneider, *Convex Bodies: The Brunn-Minkowski Theory*, Encyclopedia of Mathematics and its applications 41, Cambridge University Press, 1993.

[76] A. Schrijver, *Theory of linear and integer programming*, Wiley-Interscience, 1986.

[77] C. L. Siegel, *Lectures on the Geometry of Numbers*, Springer 1989.

[78] R. Stanley, *Enumerative Combinatorics*, Cambridge Studies in Advanced Mathematics, 62, Cambridge University Press 1999.

[79] T. Sunada, *Discrete Geometric Analysis*, Proceedings of Symposia in Pure Mathematics, 2007.

[80] G. Tardos, *Polynomial Bound for a Chip Firing Game on Graphs*, SIAM Journal on Discrete Mathematics, 1 **(3)**, 397-398, 1988.

[81] Z. B. Yossef, T. S Jayram, R. Kumar, D. Sivakumar and L. Trevisan, *Counting Distinct Elements in a Data Stream*, RANDOM: International Workshop on Randomization and Approximation Techniques in Computer Science, Lecture Notes in Computer Science, 2002,

[82] C. Zong, J. Talbot, *Sphere Packings*, Springer, 1999.