Random walks on Arakelov class groups

Boer, K. de

1. Introduction

Main concepts

In this introduction, we will treat the main concepts of this thesis in a slightly simplified and hopefully intuitive way. Though the first section roughly covers the necessary knowledge to follow this introduction, a more extensive treatment can be found in Neukirch's *Algebraic Number Theory* [NS13, Ch. 1] or Peikert's *A Decade of Lattice Cryptography* [Pei16, Sec. 2, Sec. 4]. The Arakelov class group formalism is treated nicely by Schoof [Sch08].

1.1. Number Theory

Number fields and number rings

In this thesis, the concepts of a *number field* and a *number ring* play a large role. A number field K is a finite-dimensional field extension of the rational numbers $\mathbb{Q}$, which is just a different way of saying that $K \simeq \mathbb{Q}[X]/(f(X))$ for some irreducible polynomial $f(X) \in \mathbb{Q}[X]$. The dimension of K as a $\mathbb{Q}$-vector space is called the *degree* of the number field.

Every element $\alpha \in K$ has a *minimal polynomial*, the unique monic, irreducible polynomial $m(X) \in \mathbb{Q}[X]$ satisfying $m(\alpha) = 0$. If, additionally, the minimal polynomial of α lies in $\mathbb{Z}[X]$, we call α an *integral element* of K. The integral elements in K together form a *ring*, denoted $\mathcal{O}_K$, and is named the *ring of integers* of K. Subrings of such a ring of integers of some number field K are called *number rings.*

In this introduction, we will always take the number ring to be the ring of integers $\mathcal{O}_K$ of K, for the sake of simplicity; but the ideas of this introduction apply to any other number ring $R \subseteq \mathcal{O}_K$ as well.

The Minkowski embedding

Let $K = \mathbb{Q}[X]/(f(X))$ be a number field defined by the irreducible polynomial $f(X) \in \mathbb{Q}[X]$. This polynomial $f(X)$ has $\deg(f)$ distinct roots in the complex numbers $\mathbb{C}$. This yields $\deg(f)$ different *field embeddings* $K \hookrightarrow \mathbb{C}$, respectively, by sending $\bar{X} \in K = \mathbb{Q}[X]/(f(X))$ to any of the roots of f in $\mathbb{C}$. Those are all possible field embeddings of K into $\mathbb{C}$. By concatenating these field embeddings next to each other, one gets the *Minkowski embedding* $K \to \bigoplus_{\sigma:K\to\mathbb{C}} \mathbb{C}$, $\alpha \longmapsto (\sigma(\alpha))_\sigma$. In most of the literature, the codomain of this Minkowski embedding is restricted to $K_\mathbb{R} = \{x_\sigma \in \bigoplus_{\sigma:K\to\mathbb{C}} \mathbb{C} \mid \overline{x_\sigma} = x_{\bar{\sigma}}\}$, where $\bar{\sigma}$ is the embedding $\bar{\sigma} : K \hookrightarrow \mathbb{C}$ obtained by applying first σ and then complex conjugation in $\mathbb{C}$. By componentwise addition and multiplication, $K_\mathbb{R}$ is an $\mathbb{R}$-algebra. We will see later that the ring of integers $\mathcal{O}_K$ forms a full-rank *lattice* in $K_\mathbb{R}$ under the Minkowski embedding.

Take as an example the number field $K = \mathbb{Q}[X]/(X^2 - 2)$, which has two embeddings into $\mathbb{C}$, corresponding to the zeroes $\pm\sqrt{2}$ of the polynomial $X^2 - 2$ in $\mathbb{C}$. Due to the fact that each of those actually embeds K into $\mathbb{R} \subseteq \mathbb{C}$, the (restricted) codomain $K_\mathbb{R}$ of the Minkowski embedding equals the real plane $\mathbb{R}^2$. The Minkowski embedding sends, in this case, $\bar{X} \in K$ to $(\sqrt{2}, -\sqrt{2}) \in \mathbb{R}^2$ and $1 \in K$ to $(1, 1) \in \mathbb{R}^2$ and is, by linear extension, totally determined (see Figure 1.1). Such a number field K is, by abuse of notation, often just denoted $\mathbb{Q}(\sqrt{2})$, and its ring of integers $\mathbb{Z}[\sqrt{2}]$, where $\sqrt{2}$ is used as a more understandable placeholder for $\bar{X}$. Although the ring of integers of $K = \mathbb{Q}(\alpha)$ (with α an integral element of K) equals $\mathbb{Z}[\alpha]$ in the specific examples of this introduction, this is generally not the case for other number fields.

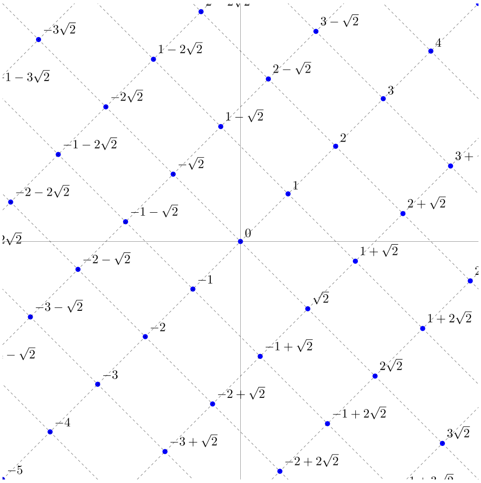

Figure 1.1.: The number ring $\mathbb{Z}[\sqrt{2}]$ visualized on the real plane, using the Minkowski embedding, sending $\sqrt{2} \mapsto (\sqrt{2}, -\sqrt{2})$ and $1 \mapsto (1, 1)$.

A slightly more intricate example concerns the number field $K = \mathbb{Q}(\sqrt[3]{2}) = \mathbb{Q}[X]/(X^3 - 2)$, which has three embeddings into $\mathbb{C}$, corresponding to the zeroes $\zeta_3^j \cdot \sqrt[3]{2}$ of the polynomial $X^3 - 2$ in $\mathbb{C}$ (where ζ_3 is a third primitive root of unity). The Minkowski embedding of K sends $\bar{X} = \sqrt[3]{2} \mapsto (\sqrt[3]{2}, \zeta_3 \cdot \sqrt[3]{2}) \in \mathbb{R} \times \mathbb{C}$ and $1 \mapsto (1, 1) \in \mathbb{R} \times \mathbb{C}$ (see Figure 1.2). Both the introduction of the reals and the absence of a third embedding is due to the restriction of the codomain of the Minkowski embedding – this third component just follows from conjugating the second component.

An appropriate metric on number fields

The Minkowski embedding $K \hookrightarrow K_{\mathbb{R}}$ yields, via the Euclidean metric on the $\mathbb{R}$-algebra $K_{\mathbb{R}}$, a *metric* on the number field K and its ring of integers $\mathcal{O}_K$. More specifically, this metric is defined via the geometric norm $\|\alpha\| := \sqrt{\sum_\sigma |\sigma(\alpha)|^2}$.

One of the advantages of this specific metric is its tight connection with the *algebraic norm* on the field K, which can be defined on $\alpha \in K$ by taking the products of the all embeddings: $\mathcal{N}(\alpha) = \prod_\sigma \sigma(\alpha)$. The algebraic and

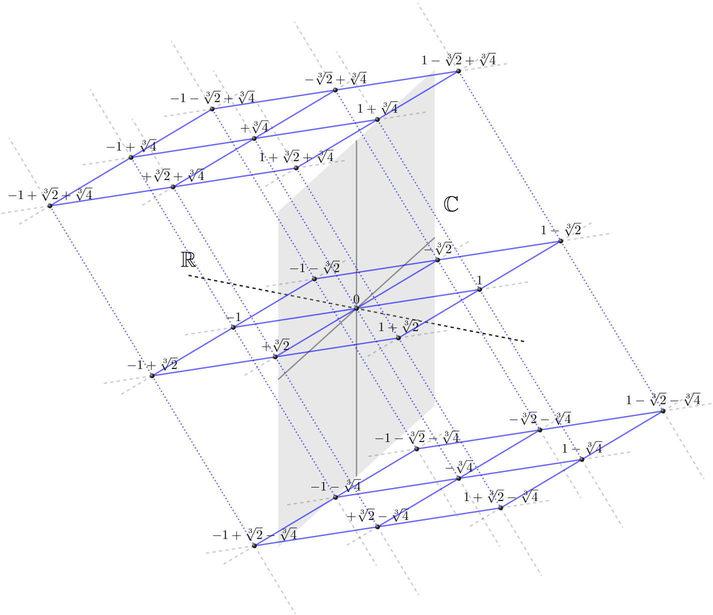
$1-\sqrt[3]{2}+\sqrt[3]{4}$
$-\sqrt[3]{2}+\sqrt[3]{4}$
$-1-\sqrt[3]{2}+\sqrt[3]{4}$
$1+\sqrt[3]{4}$
$+\sqrt[3]{4}$
$-1+\sqrt[3]{4}$
$1+\sqrt[3]{2}+\sqrt[3]{4}$
$+\sqrt[3]{2}+\sqrt[3]{4}$
$-1+\sqrt[3]{2}+\sqrt[3]{4}$
$\mathbb{C}$
$1-\sqrt[3]{2}$
$-\sqrt[3]{2}$
$\mathbb{R}$
$-1-\sqrt[3]{2}$
1
0
-1
$1+\sqrt[3]{2}$
$+\sqrt[3]{2}$
$-1+\sqrt[3]{2}$
$1-\sqrt[3]{2}-\sqrt[3]{4}$
$-\sqrt[3]{2}-\sqrt[3]{4}$
$-1-\sqrt[3]{2}-\sqrt[3]{4}$
$1-\sqrt[3]{4}$
$-\sqrt[3]{4}$
$-1-\sqrt[3]{4}$
$1+\sqrt[3]{2}-\sqrt[3]{4}$
$+\sqrt[3]{2}-\sqrt[3]{4}$
$-1+\sqrt[3]{2}-\sqrt[3]{4}$

Figure 1.2.: This picture shows the Minkowski embedding of $\mathbb{Z}[\sqrt[3]{2}]$ into $\mathbb{R} \times \mathbb{C}$.

geometric norm are related by the arithmetic-geometric mean inequality, a fact that is classically used to show the finiteness of the class group.

1.2. Ideal Lattices

Ideals

Due to the canonical geometry of the number field K, the image of the ring of integers $\mathcal{O}_K$ under the Minkowski embedding is a *discrete subgroup* in $K_\mathbb{R}$, if one only considers the additive structure of $\mathcal{O}_K$ [NS13, Ch. 1, § 4]. In other words, the ring of integers $\mathcal{O}_K$ forms a *lattice* under this embedding (see Figures 1.1 and 1.2). In fact, the same holds for any non-zero *ideal* of $\mathcal{O}_K$ in K. Recall that an ideal is a subgroup $I \subseteq \mathcal{O}_K$ of the additive group of $\mathcal{O}_K$ that is stable under multiplication with elements in $\mathcal{O}_K$, i.e., $\mathcal{O}_K \cdot I \subseteq I$.

Ideal lattices

The image of an ideal I under the Minkowski embedding is an example of an *ideal lattice*; it has the additive structure of a lattice and the ring-like structure of an ideal. An *ideal lattice* is defined as any non-zero lattice $L \subseteq K_{\mathbb{R}}$ that satisfies $\mathcal{O}_K \cdot L \subseteq L$, where the action of $\mathcal{O}_K$ happens component-wise after the Minkowski embedding (see Figure 1.3). Equivalently, considering $K_{\mathbb{R}}$ as an $\mathcal{O}_K$-algebra, ideal lattices are discrete $\mathcal{O}_K$-submodules of $K_{\mathbb{R}}$. Recall that discrete subgroups of Euclidean vector spaces correspond precisely to free $\mathbb{Z}$-modules spanned by $\mathbb{R}$-linearly independent vectors in this vector space, both called (generic) lattices.

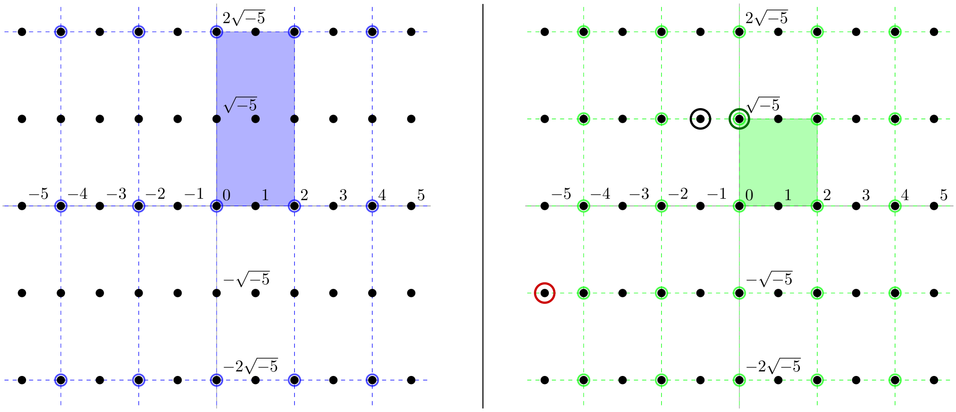

Figure 1.3.: The blue lattice is the ideal lattice $2\cdot\mathbb{Z}[\sqrt{-5}]$ in $\mathbb{Q}(\sqrt{-5})$ consisting of multiples of 2. The green lattice is an example of a lattice that is *not* an ideal lattice of $\mathbb{Q}(\sqrt{-5})$, because it is not stable under multiplication with elements of the ring of integers $\mathbb{Z}[\sqrt{-5}]$ of $\mathbb{Q}(\sqrt{-5})$. For example, $(-1+\sqrt{-5})\cdot\sqrt{-5} = -5-\sqrt{-5}$, which is the red point and does not lie in the green lattice.

It can be shown that ideal lattices L are always of the shape $L = x \cdot I$, where $I \subseteq \mathcal{O}_K$ is a non-zero ideal and $x \in K_{\mathbb{R}}^*$ (the invertible elements of $K_{\mathbb{R}}$), where the multiplication comes from the $\mathcal{O}_K$-algebra structure of $K_{\mathbb{R}}$, i.e., component-wise. In other words, ideal lattices are of the shape $L = \{(x_\sigma \cdot \sigma(\iota))_\sigma \mid \iota \in I\}$, and can be considered as ideals with a deformation. They can be stretched and squished in several coordinates by the factor $x \in K_{\mathbb{R}}^*$, see Figure 1.4.

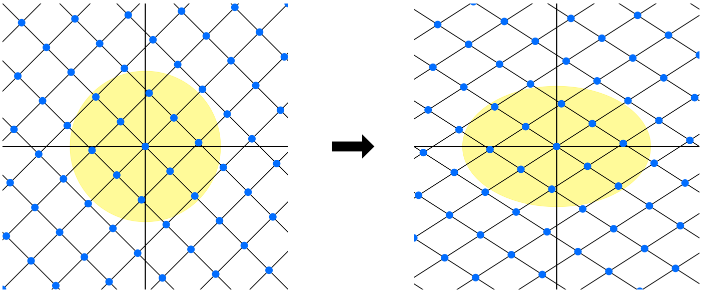

Figure 1.4.: In this two-dimensional example, the left ideal lattice is slightly stretched in the x-direction and slightly squished in the y-direction, leading to the perturbed ideal lattice on the right. The yellow circle functions as a visual aid, making the precise deformation of the lattice more explicit.

The ideal lattices within $K_{\mathbb{R}}$ form a *group* in which the multiplication is inherited from $K_{\mathbb{R}}$ and the group of (fractional) ideals; $(x \cdot I) \cdot (y \cdot J) := (x \cdot y) \cdot (I \cdot J)$, and where the unit ideal lattice is $\mathcal{O}_K \subseteq K_{\mathbb{R}}$ (under the Minkowski embedding).

In the remainder of this introduction, we will consider the group of ideal lattices *up to scaling*. This can be done by only considering ideal lattices of fixed determinant, or by constructing the equivalence relation in which two ideal lattices are equivalent if they only differ by scaling. From now on, we will refer to *this* group as the group of ideal lattices of a number field K, and we denote it by IdLat_K^0.

'Similar' ideal lattices

Next to scaling, another equivalence of ideal lattices plays a large role, one that we will call *geometrically similar* in this introductory text. Two ideal lattices $x \cdot I, y \cdot J \in \mathrm{IdLat}_K^0$ are called geometrically similar, denoted $x \cdot I \sim y \cdot J$, if there exists a $\kappa = (\kappa_\sigma)_\sigma \in K_{\mathbb{R}}$ with $|\kappa_\sigma| = 1$ for all σ, such that $\kappa \cdot x \cdot I = y \cdot J$.

The ideal lattices that are geometrically similar to the unit ideal lattice $\mathcal{O}_K$ form a subgroup called the *trivial-class ideal lattice*. In the left image of Figure 1.5 some examples of trivial-class ideal lattices are given, whose geometric similarity with $\mathcal{O}_K$ can be verified visually.

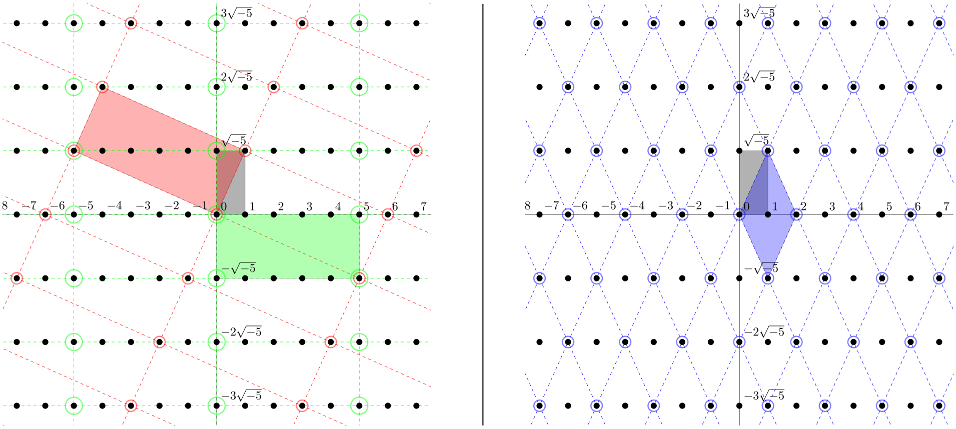

Figure 1.5.: On the left are three trivial-class ideal lattices, namely $\mathbb{Z}[\sqrt{-5}]$, $\sqrt{-5} \cdot \mathbb{Z}[\sqrt{-5}]$ and $(1 + \sqrt{-5}) \cdot \mathbb{Z}[\sqrt{-5}]$. By observing the rectangular shapes enclosed by the lattice points, one indeed observes that these three lattices are equal, up to scaling and rotation, and thus geometrically similar. In the right image one can see the blue ideal lattice, which is the smallest ideal lattice in $\mathbb{Z}[\sqrt{-5}]$ containing both 2 and $\sqrt{-5}$. As the shape of this lattice is a diamond instead of a rectangle, it cannot be a trivial-class ideal lattice.

1.3. Arakelov Class Groups

Looking again at the ideal lattices of $\mathbb{Q}(\sqrt{-5})$ in Figure 1.5, we can distinguish two shapes of ideal lattices; a rectangle with proportion $\sqrt{5} : 1$, and a diamond with height $\sqrt{5}$ and width 2. A reasonable question to ask is: *do all ideal lattices in* $\mathbb{Q}(\sqrt{-5})$*, up to scaling and geometric similarity, fall into one of these two shapes?* The answer turns out to be *yes*; this is closely related to the fact that $\mathbb{Q}(\sqrt{-5})$ is a complex quadratic number field with class number two.

Summarizing, the ideal lattices in $\mathbb{Q}(\sqrt{-5})$ fall into two *classes*, the 'rectangle' class ▬ and the 'diamond' class ♦. This categorization of the ideal lattices of $\mathbb{Q}(\sqrt{-5})$ is described by the *Arakelov class group* of $\mathbb{Q}(\sqrt{-5})$, which we denote by $\mathrm{Pic}^0_{\mathbb{Q}(\sqrt{-5})}$. In other words,

$$\mathrm{Pic}^0_{\mathbb{Q}[\sqrt{-5}]} = \{\text{🟥}, \text{🔷}\}.$$

This categorization of ideal lattices can be done for any number field; in fact, we have the following definition of the Arakelov class group.

> The Arakelov class group of a number field is the *group of geometric similarity classes of ideal lattices* of that number field.
> Symbolically,
>
> $$\mathrm{Pic}^0_K := \mathrm{IdLat}^0_K \,/\!\sim,$$
>
> where $\sim$ is the equivalence relation of being geometrically similar.

The shapes of the ideal lattices of $\mathbb{Q}(\sqrt{-5})$ fall into *two* classes, in other words, $|\,\mathrm{Pic}^0_{\mathbb{Q}(\sqrt{-5})}\,| = 2$, a *finite* number. The Arakelov class group being finite only happens in imaginary quadratic number fields and the rationals $\mathbb{Q}$, for which can be shown that it is canonically isomorphic to the ideal class group.

In all other number fields the Arakelov class group is an infinite (but compact) abelian group. A way of visualizing this is by imagining a spectrum of lattice shapes; so, for example, not only diamond-shaped or rectangle-shaped, but also everything in between, see Figure 1.6.

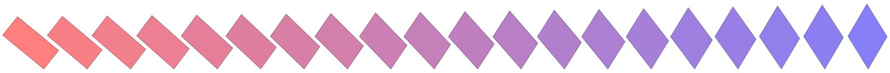

Figure 1.6.: In most number fields, the Arakelov class group is infinite. The ideal lattices have an infinite variety of shapes. For example, one could imagine that these shapes shift seamlessly from 🟥 to 🔷.

Infinite Arakelov class groups

We will now consider an example of a number field whose Arakelov class group is infinite, namely that of $\mathbb{Q}(\sqrt{3})$, with ring of integers $\mathbb{Z}[\sqrt{3}]$. To show that there is a larger variety of ideal lattices here, we refer to Figure 1.7 for three examples of shapes of ideal lattices in $\mathbb{Z}[\sqrt{3}]$.

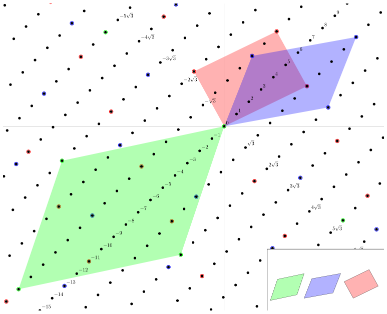

Figure 1.7.: In this picture we examine three different shapes of ideal lattices of the number field $\mathbb{Q}(\sqrt{3})$. The variety of shapes suggests that there is a spectrum of different ideal lattice shapes.

Indeed, a continuous spectrum of ideal lattice shapes happen to occur in $\mathbb{Q}(\sqrt{3})$, slightly similar to Figure 1.6. Furthermore, this spectrum of ideal lattice shapes can be exactly found by stretching the shape in the x-direction and shrinking the same amount in the y-direction (and vice versa, see Figure 5.2). The deformation of ideal lattices in this way is possible because the field $\mathbb{Q}(\sqrt{3})$ has *two* independent (real) embeddings into $\mathbb{C}$, as opposed to $\mathbb{Q}(\sqrt{-5})$, which has only one independent (complex) embedding[1]. Note that the product of the deformations in the x and y-direction is required to be 1, in order to keep the the determinant of the ideal lattice fixed.

Changing an ideal lattices shape this way, something peculiar occurs eventually: at a certain point of deforming the lattice shape, one arrives at a different shape, *but representing the same lattice*; an example of this phenomenon can be seen in Figure 1.8. As a result, the Arakelov class group of $\mathbb{Q}(\sqrt{3})$ has a circular nature, and is in fact isomorphic to the *circle group* S^1, see Figure 1.9.

More explicitly, the ideal lattice group has the following parametrization for

[1] Technically, imaginary quadratic number fields have two embeddings into the complex space, but they are dependent in the way that one is the complex conjugate of the other.

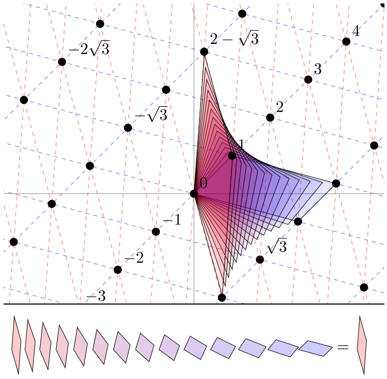

Figure 1.8.: We consider here ideal lattices of the number field $\mathbb{Q}(\sqrt{3})$. By stretching the red shape into the x-direction and shrinking the same amount in the y-direction, one obtains a variety of shapes. Eventually, one arrives at the blue shape, which represents the same ideal lattice as that of the red shape. In fact, this full spectrum of ideal lattices in $\mathbb{Q}(\sqrt{3})$ is precisely obtained by multiplying $\mathbb{Z}[\sqrt{3}]$ by (e^t, e^{-t}) for $t \in [0, \log(2+\sqrt{3})]$, where $2+\sqrt{3}$ is the fundamental unit of $\mathbb{Z}[\sqrt{3}]$.

$t \in \mathbb{R}$,

$$\mathrm{IdLat}^0_{\mathbb{Q}(\sqrt{3})} = \{(e^t, e^{-t}) \cdot \mathbb{Z}[\sqrt{3}] \subseteq K_{\mathbb{R}} \mid t \in \mathbb{R}\}.$$

The ring of integers $\mathbb{Z}[\sqrt{3}]$ has the element $2+\sqrt{3} = (2-\sqrt{3})^{-1}$ as a *fundamental unit*, and therefore, taking $t = \log(2+\sqrt{3})$, we have

$$(e^t, e^{-t}) \cdot \mathbb{Z}[\sqrt{3}] = (2+\sqrt{3}, 2-\sqrt{3}) \cdot \mathbb{Z}[\sqrt{3}] = \mathbb{Z}[\sqrt{3}] = (e^0, e^0) \cdot \mathbb{Z}[\sqrt{3}].$$

As a result, the Arakelov class group of $\mathbb{Q}(\sqrt{3})$ is, via the above parametrization, isomorphic to $\mathbb{R}/\log(2+\sqrt{3}) \cdot \mathbb{Z}$, a circle group. So, the Arakelov class group $\mathrm{Pic}^0_K \simeq \mathbb{R}/\log(2+\sqrt{3}) \cdot \mathbb{Z}$ of $\mathbb{Q}(\sqrt{3})$, has *volume* (length) $\log(2+\sqrt{3})$, which is exactly the *regulator* R of the number field $\mathbb{Q}(\sqrt{3})$.

This is not a coincidence. In this specific case, because $\mathbb{Q}(\sqrt{3})$ has class number one, the Arakelov class group is canonically isomorphic to the *quotient group* $H/\mathrm{Log}(\mathcal{O}_K^\times)$ *of the hyperplane* $H = \mathrm{span}(\mathrm{Log}(\mathcal{O}_K^\times))$ *and the logarithmic unit lattice* $\mathrm{Log}(\mathcal{O}_K^\times)$ that arises in Dirichlet's unit theorem.

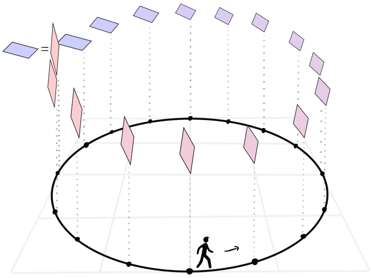

Figure 1.9.: By deforming an initial ideal lattice in $\mathbb{Q}(\sqrt{3})$ appropriately, one eventually arrives at the same ideal lattice. This yields a circular pattern; as a result, the Arakelov class group of $\mathbb{Q}(\sqrt{3})$ is isomorphic to the circle group S^1.

General Arakelov class groups

In the previous text, we saw two examples of an Arakelov class group. One of an imaginary quadratic number field $\mathbb{Q}(\sqrt{-5})$, which was a finite group isomorphic to the class group, and one of a real quadratic number field $\mathbb{Q}(\sqrt{3})$ which was isomorphic to a circle with the volume equal to the regulator.

So, in one case the Arakelov class group seems tightly related to the *class group*, whereas in another case it seems related to the *unit group*. In reality, it is related to *both*: it is a 'combination' of both the class group $\mathrm{Cl}(K)$ and the *logarithmic unit torus* $T = H/\operatorname{Log}(\mathcal{O}_K^\times)$, the quotient group of the hyperplane $H = \operatorname{span}(\operatorname{Log}(\mathcal{O}_K^\times))$ and the logarithmic unit lattice $\operatorname{Log}(\mathcal{O}_K^\times)$. Here, $\operatorname{Log}(\eta) := (\log|\sigma(\eta)|)_\sigma$ for $\eta \in \mathcal{O}_K^\times$ is the *logarithmic map*, defined by taking the component-wise logarithm of the absolute values of the Minkowski embedding. This turns the multiplicative group of units $\mathcal{O}_K^\times$ into a lattice $\operatorname{Log}(\mathcal{O}_K^\times)$, of which the hyperplane H is the linear span.

More precisely, the Arakelov class group fits in an exact sequence where the outer groups are the class group $\mathrm{Cl}(K)$ and the logarithmic unit torus $T = H/\operatorname{Log}(\mathcal{O}_K^\times)$.

$$0 \to H/\operatorname{Log}(\mathcal{O}_K^\times) \to \operatorname{Pic}_K^0 \to \mathrm{Cl}(K) \to 0.$$

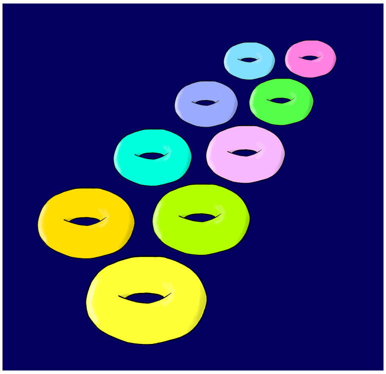

Figure 1.10.: The Arakelov class group of a number field K consists of a union of finitely many hypertori.

The specific cases of $\mathbb{Q}(\sqrt{-5})$ and $\mathbb{Q}(\sqrt{3})$ can now reasonably be explained. In the imaginary quadratic case of $\mathbb{Q}(\sqrt{-5})$ the logarithmic unit torus $T = H/\operatorname{Log}(\mathcal{O}_K^\times)$ consists of a single point (due to the unit group being finite), which makes the Arakelov class group isomorphic to the class group. In the real quadratic case $\mathbb{Q}(\sqrt{3})$, however, the class group is trivial instead, so that the Arakelov class group is isomorphic to the logarithmic unit torus $T = H/\operatorname{Log}(\mathcal{O}_K^\times) \simeq \mathbb{R}/\log(2+\sqrt{3})\mathbb{Z}$, i.e., $\mathbb{R}$ quotiented out by the free group generated by the logarithm of the fundamental unit of $\mathbb{Q}(\sqrt{3})$; this is a circle group.

In the most general case, the logarithmic unit torus $T = H/\operatorname{Log}(\mathcal{O}_K^\times)$ is a *hypertorus* and the class group is a finite abelian group. This leads to the following topological description of the Arakelov class group.

The Arakelov class group of a number field K consists of a union of finitely many hypertori. The number of tori is equal to the class number of K and all tori are isomorphic to the logarithmic unit torus $T = H/\operatorname{Log}(\mathcal{O}_K^\times)$, thus having a volume equal to the regulator of K.

Summarizing, a *point* on a torus in the Arakelov class group corresponds to an *ideal lattice* (more precisely, a class of same-shaped ideal lattices) in the number field. Moving the point on the torus a little corresponds to slightly disturbing the shape of the lattice, exactly like in the circle of Figure 1.9 (see also Figure 1.11).

If the corresponding points of two lattices lie on the same torus (in the Arakelov class group), they can be transformed into each other by means of stretching and shrinking appropriately. If, on the other hand, these points lie on *different* tori of the Arakelov class group, they *can not* be transformed in one another, see Figure 1.11.

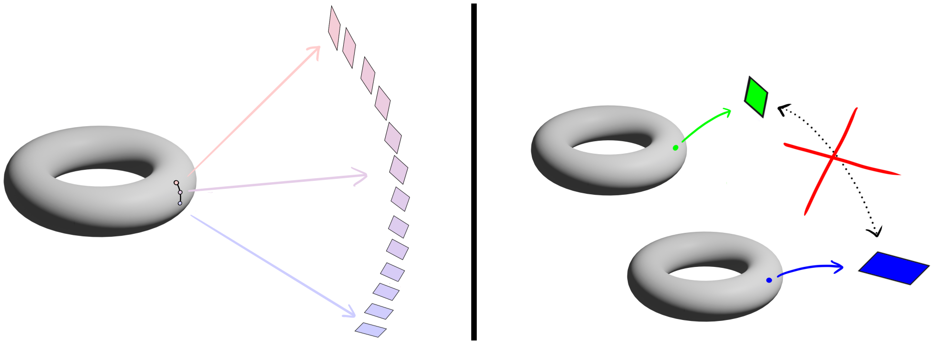

Figure 1.11.: Any two lattices corresponding to points on the *same* torus, can be transformed into each other (left). If two lattices correspond to points on *different* tori, they cannot be transformed into each other (right).

1.4. Random Walks on the Arakelov Class Group

The main theorem of this thesis involves *random walks* on the Arakelov class group, a specific algorithm that allows to move randomly.

What is a random walk?

An intuitive way of thinking about a random walk is by picturing an ant on a plane, where the ant gets no external stimuli. This ant will move in random directions with quite an irregular path, see the left-most picture of Figure 1.12.

Due to the random behavior of the ant, we do not know its precise future movements. So, in order to predict the ant's future position, we have to resort to using stochastics. The probability distribution that describes the possible end points of the ant after a certain given time is called the *random walk distribution*, and will, on the real plane, take the shape of a Gaussian distribution (see the right-most picture of Figure 1.12).

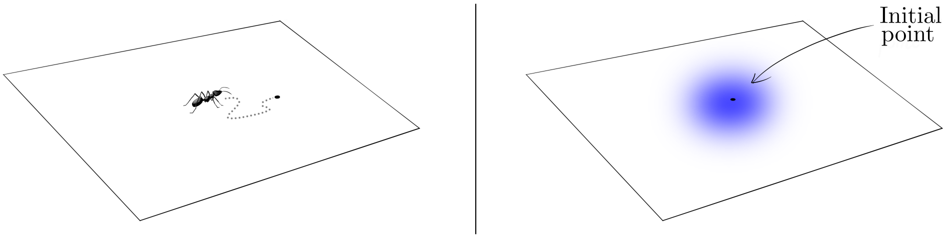

Figure 1.12.: An ant on a plane will, without external stimuli, follow an irregular path, as in the left-most image. This is can be regarded as an intuitive interpretation of a random walk. The probability distribution arising from this statistical behavior is called a *random walk distribution* and is visualized in the right-most image.

One can actually define a random walk on *any* reasonable surface (or even in the three-dimensional or higher-dimensional space, by imagining a confused fly). The most relevant surface for our purposes is the *hypertorus*, because that is what an Arakelov class group consists of.

In random walks on hypertori, something peculiar occurs whenever the deviation of the Gaussian gets large. Namely, at a certain deviation the Gaussian distribution 'folds round' the entire hypertorus, and is evenly spread out everywhere; this concept is known as *smoothing* in the theory of lattices. So, this Gaussian random walk distribution on a torus, with increasing deviation, tends to a uniform distribution, see Figure 1.13.

How to randomly walk on the Arakelov class group?

The Arakelov class group consists of finitely many hypertori. Each point on one of these tori corresponds to a lattice geometric-similarity class, and

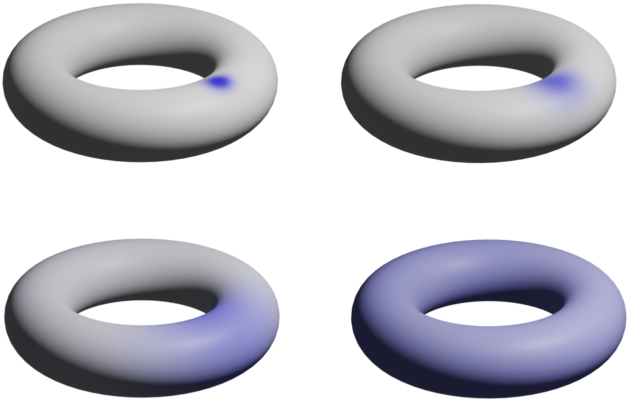

Figure 1.13.: As the deviation of the Gaussian distribution increases, the distribution 'folds around' the torus more and more. From a certain value of the deviation, the distribution is very close to a uniform distribution.

deforming this lattice allows to move around on one torus, see Figure 1.11. However, in order to obtain a reasonable covering random walk on an Arakelov class group we need to be able to *jump* from one torus to the other as well.

Before unveiling yet how we actually achieve such a jump in terms of lattices, we define the two allowed moves in a random walk on the Arakelov class group.

- 'Crawling', that is, (slowly) moving on one single torus.
- 'Jumping', that is, instantaneously teleport (as it were) to a certain distant point either on a different torus, or on the same torus.

Because of these two movements, an ant is not anymore the appropriate insect to keep in mind for intuition. Instead, we might want to think of a *grasshopper*, see Figure 1.14.

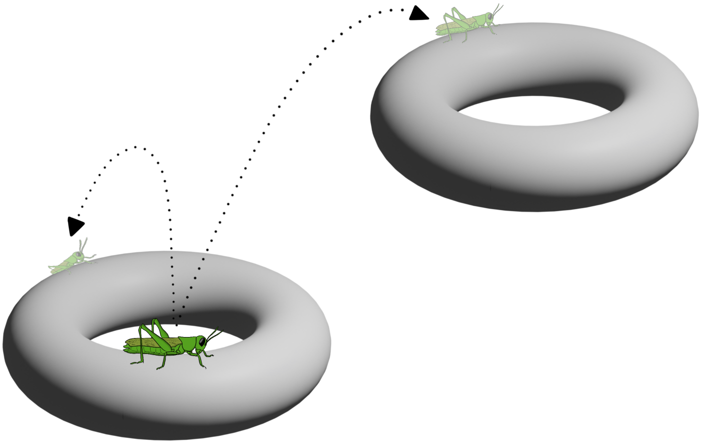

Figure 1.14.: Due to the disconnected nature of the Arakelov class group, as it consists of several separated tori, we also need ‘jumps’ in our random walk, next to ‘crawls’. For intuition it is then more appropriate to have a grasshopper in mind. The grasshopper does not need to land on a different torus per se, but can also jump to a distant place on the same torus.

‘Crawling’ by a Gaussian deformation

In terms of ideal lattices, ‘crawling’ happens by multiplying the input ideal by a random log-normal deformation $x \in K_{\mathbb{R}}$ satisfying $\prod_\sigma x_\sigma = 1$, in order to keep the determinant of the ideal lattice the same.

More precisely, we pick a Gaussian vector $(g_\sigma)_\sigma$ in which each entry is a zero-centered Gaussian with deviation s, subject to the requirement $\sum_\sigma g_\sigma = 0$. Putting $x_\sigma = e^{g_\sigma}$ yields the correct log-normal distribution on $K_{\mathbb{R}}$.

‘Jumping’ by multiplying with prime ideals

In terms of ideal lattices, such a jump from one torus to another happens by *multiplying* the initial ideal lattice by a (non-zero) *prime ideal*. More specifically, denoting $\mathfrak{p} \subseteq \mathcal{O}_K$ for a prime ideal of $\mathcal{O}_K$, the operation $x \cdot \mathfrak{a} \longmapsto x \cdot (\mathfrak{p} \cdot \mathfrak{a})$ yields a jump in the Arakelov class group[2].

Geometrically, multiplying an ideal lattice $L = x \cdot \mathfrak{a}$ by a prime ideal of $\mathcal{O}_K$ corresponds to taking a *prime sub ideal lattice* $x \cdot (\mathfrak{p} \cdot \mathfrak{a}) \subseteq x \cdot \mathfrak{a}$, that is a sub ideal lattice $P \subseteq x \cdot \mathfrak{a}$ for which no proper ideal lattice lies in between. In other words, for a prime sub ideal lattice $P \subseteq x \cdot \mathfrak{a}$ there are no ideal lattices L such that $P \subsetneq L \subsetneq x \cdot \mathfrak{a}$ (see Figure 1.15).

As we would like the jumps to other tori to be random, aimless like a grasshopper, a *probabilistic* element is added. Starting from a certain initial ideal lattice $x \cdot \mathfrak{a}$ (corresponding to a point on the Arakelov class group), we uniformly random pick a prime ideal $\mathfrak{p} \subseteq \mathcal{O}_K$ among all prime ideals with norm bounded by some bound B, and switch to the lattice $x \cdot (\mathfrak{p} \cdot \mathfrak{a})$. This procedure of multiplying by a random prime can be repeated as often as we want; we denote with N the total number of these ‘jumps’. A toy example with two jumps (so $N = 2$) is depicted in Figure 1.16.

[2] To be completely precise, it would be more correct to write $x \cdot \mathfrak{a} \longmapsto (x \cdot \mathcal{N}(\mathfrak{p})^{-1/n}) \cdot (\mathfrak{p} \cdot \mathfrak{a})$, where the norm $\mathcal{N}(\mathfrak{p})$ of $\mathfrak{p}$ is involved in order to keep the determinant fixed.

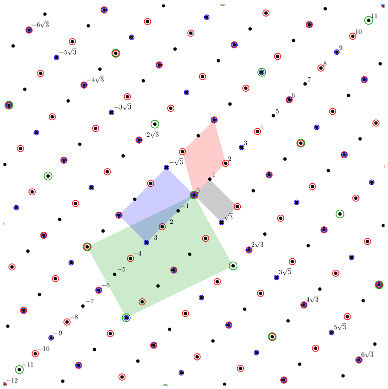

Figure 1.15.: The red, blue and green ideal lattices are all *prime* sub ideal lattices of the gray (base) ideal lattice, because the shapes of the respective ideal lattices are 2, 3 and 11 (prime numbers) times larger than the surface of the gray ideal lattice.

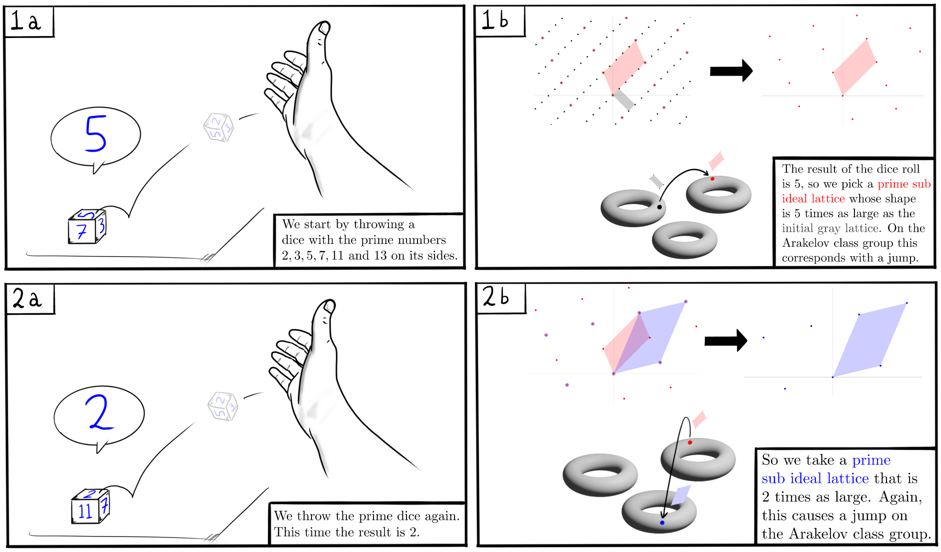

We start by throwing a dice with the prime numbers 2, 3, 5, 7, 11 and 13 on its sides.

The result of the dice roll is 5, so we pick a prime sub ideal lattice whose shape is 5 times as large as the initial gray lattice. On the Arakelov class group this corresponds with a jump.

We throw the prime dice again. This time the result is 2.

So we take a prime sub ideal lattice that is 2 times as large. Again, this causes a jump on the Arakelov class group.

Figure 1.16.: This picture shows two repetitions ($N = 2$) of a random jump. In more realistic cases, both the number of primes and the number of jumps are larger. Note that at each jump, the ideal lattice gets sparser, or, equivalently, its shape gets larger.

Description of the full random walk on the Arakelov class group

We now give the final description of our definition of a random walk on the Arakelov class group, involving three parameters: N, B and s (see Figure 1.18). Here N is the number of consecutive jumps on the Arakelov group, as well as the number of prime ideals one multiplies the input ideal with. The number B is the bound on the norms of these prime ideals and equals (up to a logarithmic factor) the number of primes one can randomly pick from in each jump. These two parameters N and B concern the 'discrete part' of the random walk. The 'continuous part' of the random walk on the other hand is determined by the deviation s of the log normal distribution of the random deformation.

A random walk on the Arakelov class group, starting from an ideal lattice $x \cdot \mathfrak{a}$, consists of two separate parts. The first part involves N random 'jumps', carried out by multiplying the ideal lattice by N random primes with bounded norm B, yielding the operation $x \cdot \mathfrak{a} \longmapsto x \cdot (\prod_{j=1}^{N} \mathfrak{p}_j)\mathfrak{a}$.

The second part, that comes after, involves a random log-normally distributed crawl $y \in K_{\mathbb{R}}$ of deviation s, which is executed by slightly deforming the lattice $x \cdot (\prod_{j=1}^{N} \mathfrak{p}_j \mathfrak{a})$ resulting from the jumps:

$$x \cdot (\prod_{j=1}^{N} \mathfrak{p}_j \mathfrak{a}) \longmapsto (y \cdot x) \cdot (\prod_{j=1}^{N} \mathfrak{p}_j \mathfrak{a}).$$

The random walk process is depicted in Figure 1.17.

1.5. The Random Walk Theorem for Arakelov Class Groups

We are now almost ready to phrase the main result of this thesis. Recalling the framework of the random walk: we tried before to predict the position of an ant walking on a torus for a certain time, only knowing its initial

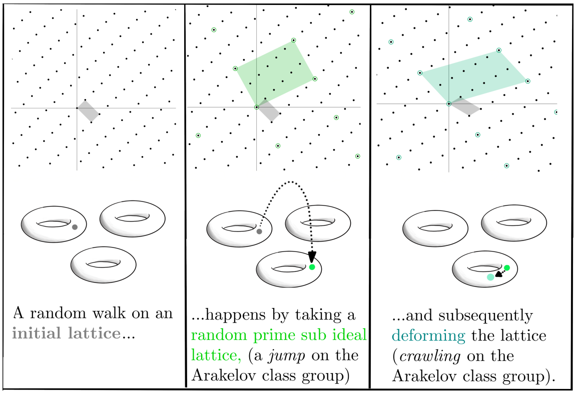
A random walk on an initial lattice...
...happens by taking a random prime sub ideal lattice, (a *jump* on the Arakelov class group)
...and subsequently deforming the lattice (*crawling* on the Arakelov class group).

Figure 1.17.: A concrete realization of the random walk procedure on an ideal lattice with a single jump. The prime sub ideal lattice is chosen at random, as well as the deformation (by sampling a Gaussian distribution).

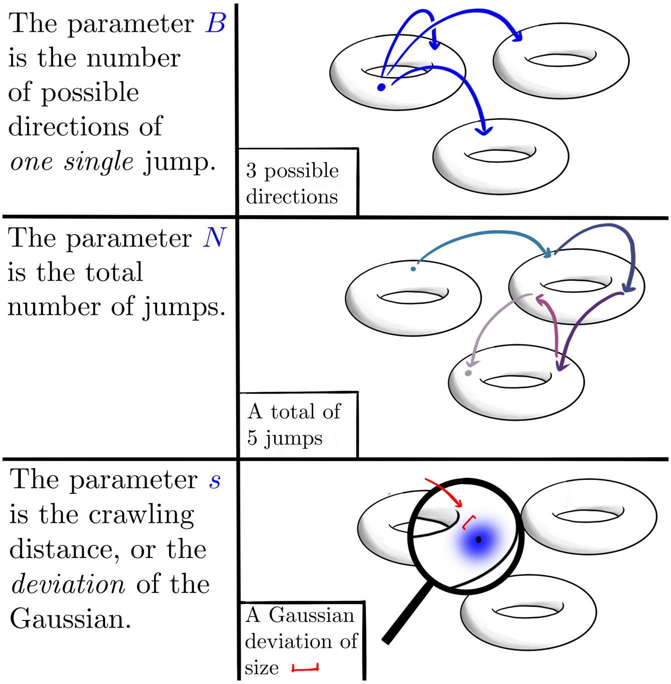
The parameter B is the number of possible directions of *one single* jump.
3 possible directions
The parameter N is the total number of jumps.
A total of 5 jumps
The parameter s is the crawling distance, or the *deviation* of the Gaussian.
A Gaussian deviation of size

Figure 1.18.: An explanation of the random walk's parameters.

position. The current situation is not much different; we now try to predict the position of a *grasshopper* on multiple tori (the Arakelov class group), only knowing its initial position, the number of jumps N, the number B of primes[3] to sample from, and the deviation s of the crawl.

As we saw in Figure 1.13, an ant's crawl of a large enough deviation 'folds around the torus' and therefore leads to a uniform distribution. Something very similar happens with the grasshopper and the Arakelov class group consisting of multiple tori. For appropriately many jumps N, appropriately many primes B and an appropriately large deviation s, the random walk distribution on the Arakelov class group is also close to the uniformly random distribution.

Intuitively, the more jumps (i.e., larger N) happen in the random walk, the less crawling (i.e., smaller s) is needed to cover[4] the Arakelov class group. The converse is also true; in the case of few jumps, more crawling is required to cover all tori, see Figure 1.19.

In the following informal geometric volume-covering argument we show a necessary condition on the parameters in order to cover the Arakelov class group fully with a random walk. In fact, if one assumes the extended Riemann hypothesis, we can show that that this necessary condition is also almost sufficient – only a slightly more larger parameter choice is sufficient to have a covering random walk.

Volume covering argument

Assume for the moment that the multiple Gaussians caused by the crawling do *not overlap at all*. Then, the total volume covered by the random walk distribution equals $\binom{B}{N} \cdot s^d$, namely, each of the $\binom{B}{N}$ possible final jump points

[3] Formally, this was the bound B on the norms of the primes; the number of primes with norm bounded by B equals $B/\log B$ so it does not much harm to identify the number of primes with B.

[4] 'Cover', here, is used in an informal sense, and not in the (formal) topological sense. In the informal geometric argument that follows, a point on the Arakelov class group is 'covered' if the random walk distribution has a non-negligible density value there.

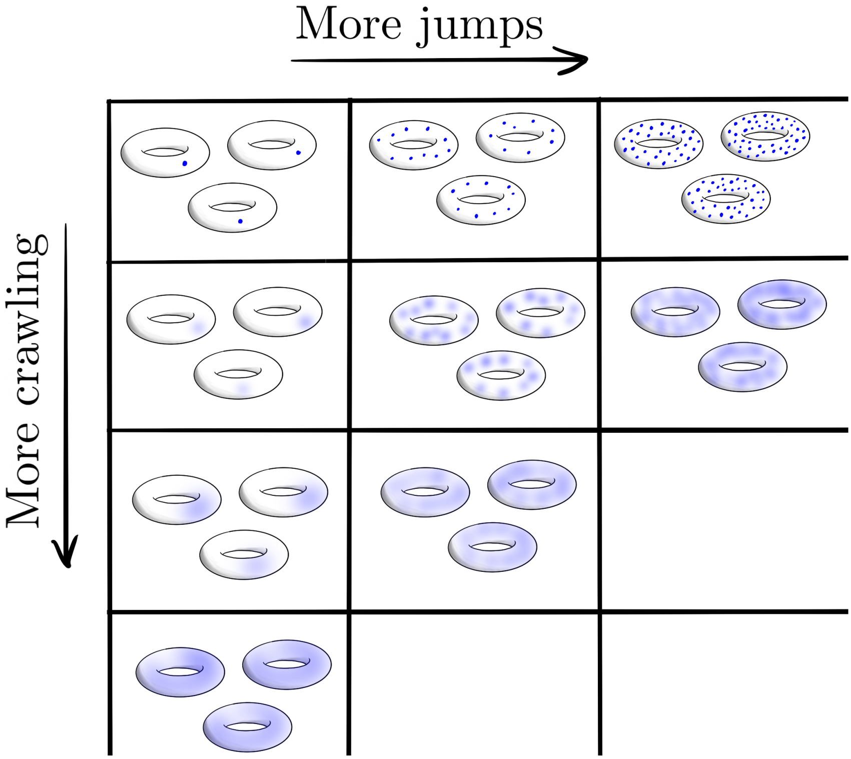

Figure 1.19.: The more jumps happen in the random walk, the less crawling is needed in order to cover the entire Arakelov class group.

have a covering of about s^d due to the crawling. Here, d is the dimension of the Arakelov class group of K which is equal to the rank of the unit group of K.

This means that for the random walk distribution to be uniform, i.e. covering everything equally, it *must* cover the entire volume of the Arakelov class group. In particular, the volume $\binom{B}{N} \cdot s^d$ covered by the random walk distribution (assuming no overlap) must exceed the volume of the Arakelov class group.

For the random walk distribution on the Arakelov class group Pic_K^0 to be uniform, the estimated volume coverage $\binom{B}{N}s^d$ of the random walk is *required* to exceed the volume $\mathrm{vol}(\mathrm{Pic}_K^0)$ of the Arakelov class group, that is,

$$\binom{B}{N} \cdot s^d \geq \mathrm{vol}(\mathrm{Pic}_K^0). \tag{1.1}$$

The assumption that the Gaussians of the random walk do not overlap at all is not a realistic one, because there will always be *some* overlap, especially whenever the covered volume almost exceeds $\mathrm{vol}(\mathrm{Pic}_K^0)$. The volume argument still holds if the overlap is just not too severe, which exactly happens if the end points of the jumps are *reasonably equidistributed.* Such equidistribution of prime ideals is often tackled by assuming some extended form of the *Riemann hypothesis*, on which we will elaborate later.

In fact, if we indeed assume this extended form of the Riemann hypothesis, we can deduce that the number of jumps N, the number of jump directions B (number of prime ideals) and the deviation s only need to be *slightly* larger than required in Equation (1.1), in order for the random walk to be uniform on the Arakelov class group. This means that the result is very near what one optimally would expect. The precise, non-simplified analogue of this statement, which is the main theorem of this thesis, is spelled out in Theorem 4.3.

The extended Riemann hypothesis

The Riemann hypothesis is at its very essence an assumption on the regularity or evenness of the prime numbers among the rest of the numbers. This assumption is often used in mathematics, mostly to prove efficiency of certain algorithms involving prime numbers.

In this thesis, we assume an extended form of this Riemann hypothesis, because we are not dealing with prime numbers, but with prime ideals. The formal statement of the Extended Riemann Hypothesis in this thesis is that

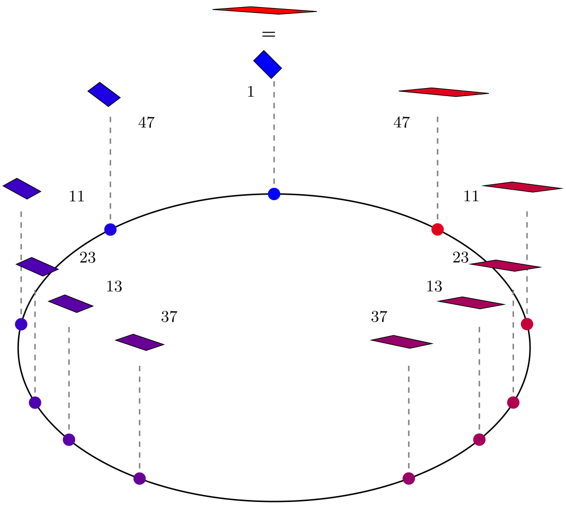

Figure 1.20.: The circle-shaped Arakelov class group of $\mathbb{Q}(\sqrt{3})$ with the positions of the first few prime ideals $\mathfrak{p}$ and their associated shapes. Already in this 'small' example, there is a reasonable equidistribution of these prime ideals on the Arakelov class group.

it assumes that all zeroes in the critical strip of Hecke L-functions of number fields lie on the $\Re(z) = 1/2$ line, see Definition 2.10. The impact is that prime ideals of a number ring lie quite equidistributed on the Arakelov class group, see Figure 1.20. For the volume covering argument of this section to be near-optimal, such equidistribution of prime ideals is of fundamental importance, which suggests the necessity of this particular form of the Riemann hypothesis. In the actual proof of the random walk theorem, this Extended Riemann hypothesis indeed seems to be indispensable (see the proof of Theorem 4.3 in Chapter 4).

1.6. A Worst-case to Average-case Reduction

Introduction

A reason why random walks on Arakelov class groups are interesting, is because of their applications. In this section we will explain one of these applications, which concerns a worst-case to average-case connection for finding short vectors in ideal lattices.

The shortest vector problem

A computational problem that plays are large role in cryptography, is called the ‘shortest vector problem’. The associated computational question is to find a short non-zero point (vector) in a given lattice. Short, here, means that the lattice point needs to be *close* to the origin, but not the origin itself, see Figure 1.21.

More precisely, for a given lattice L, the *r-approximate shortest vector problem* (approx-SVP) is the problem of finding a non-zero lattice point $\ell \in L$ that satisfies $\|\ell\| < r$. When only lattices of fixed determinant are considered, this is named the *Hermite* approximate shortest vector problem.

Though this computational problem looks rather easy in two dimensions, it becomes more and more hard with increasing dimension. It is believed that this is true not only for classical computers, but also for quantum computers.

This is one of the reasons why this particular computational problem lies at the foundation of many post-quantum cryptographic protocols (which require an underlying ‘hard’ problem). Such cryptographic protocols based on the shortest vector problem derive their general security from the hardness of this particular problem. Because of this reason, it is of fundamental importance to analyze this hardness.

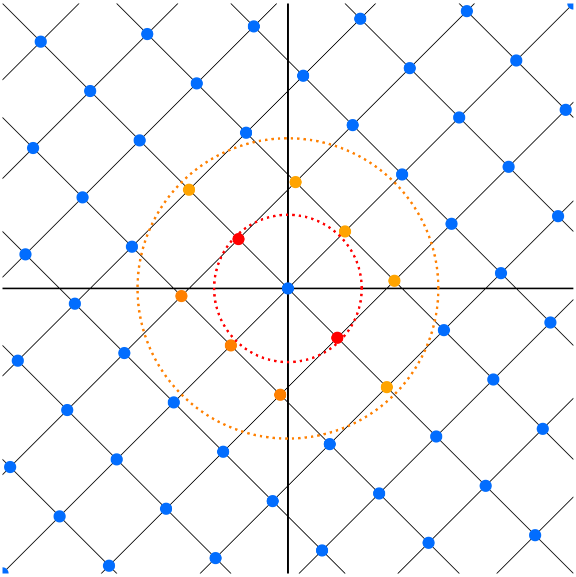

Figure 1.21.: The *shortest vector problem* asks to find a short vector in the lattice, which means that it is close to the origin, but not the origin point itself. The red points are the *shortest* lattice elements. In most cases, though, just short vectors, like the orange points, are also good. Concretely, whether a lattice point is short or not is often decided by whether the lattice point lies in a circle with predescribed radius r or not.

The Shortest Vector Problem on ideal lattices

In this thesis, we focus on the hardness of the shortest vector problem in *ideal lattices*. Ideal lattices are a special subclass of general lattices that arise from number fields. Due to this fact, as we saw in an earlier section, ideal lattices (of a fixed number field) can be assembled into geometrically equivalent classes, yielding the *Arakelov class group*. Because for two geometrically equivalent lattices it is believed to be precisely equally hard to find short vectors in, this Arakelov class group is appropriate to consider.

In this thesis, we study the hardness of finding short vectors in ideal lattices, in a *relative sense*. Concretely, one of the research questions of this thesis can be phrased as follows: is finding short vectors about equally hard for all classes of ideal lattices (case A), or are there ideal lattices in which short vectors are significantly harder to find (case B)? By giving the 'hard' ideal lattice classes a red color, and the 'easy' ideal lattice classes a green color on the Arakelov class group, these two cases are portrayed in Figure 1.22.

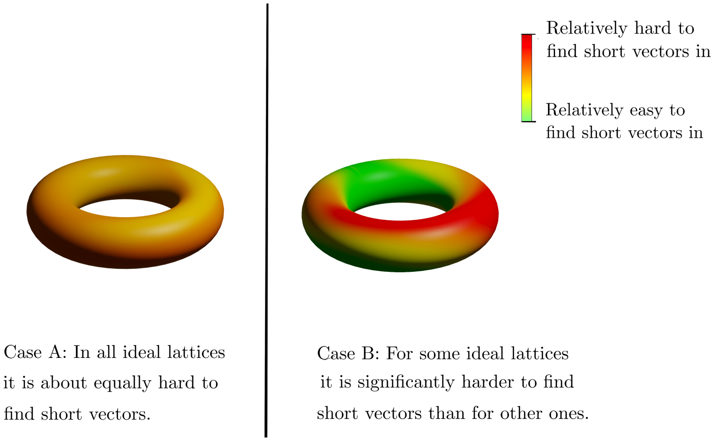
Relatively hard to find short vectors in
Relatively easy to find short vectors in

Case A: In all ideal lattices it is about equally hard to find short vectors.

Case B: For some ideal lattices it is significantly harder to find short vectors than for other ones.

Figure 1.22.: Is it for all ideal lattices on the Arakelov class group about equally hard to find short vectors in (case A) or not (case B)? Note that we just pictured one single torus for the Arakelov class group, for simplicity.

Though the full answer to this research question on relative hardness is slightly more subtle, and will be elaborated on in the next section, the simplified answer is short.

In all ideal lattices associated with a fixed number field it is about equally hard to find short vectors. In other words, Case A of Figure 1.22 is quite an accurate rendition of reality.

Argument for the evenness of this hardness on the Arakelov class group, using random walks

To give an argument *why* all ideal lattice classes in the Arakelov class group are about equally hard to find short vectors in, one can use the random walk theorem on Arakelov class groups. This argument is based on the following important observation, which is, for sake of brevity, specialized to the case of cyclotomic fields.

For cyclotomic fields, considering ideal lattices of fixed determinant, finding a lattice vector of length r in the lattice at the *end of the random walk* allows to find a short element of length $r \cdot \sqrt{n}$ in the *initial lattice*, by ‘undoing’ the random walk on the found short element, see Figure 1.23.

This observation rules out the existence of an ideal lattice in which it is (compared to other ideal lattices) extraordinarily hard to find short vectors in (such a hard lattice would be an intense red point on the Arakelov class group in Case B of Figure 1.22). Namely, by the above observation (and Figure 1.23), finding short vectors in the end lattice and in the initial lattice or a random walk is somehow very related. Therefore, finding a short vector in the fixed initial lattice cannot be so much harder than finding short vectors in a random ‘average’ lattice. Summarizing, there cannot be much variation in hardness of finding short vectors in ideal lattices, as visualized in Case A in Figure 1.22.

Note that the random walk on the Arakelov class group reduces the shortest vector problem on an initial lattice to a shortest vector problem on the end lattice with a *harder approximation factor*, since it is smaller. So, the portrayal in Figure 1.22 is not completely accurate, since we leave out this

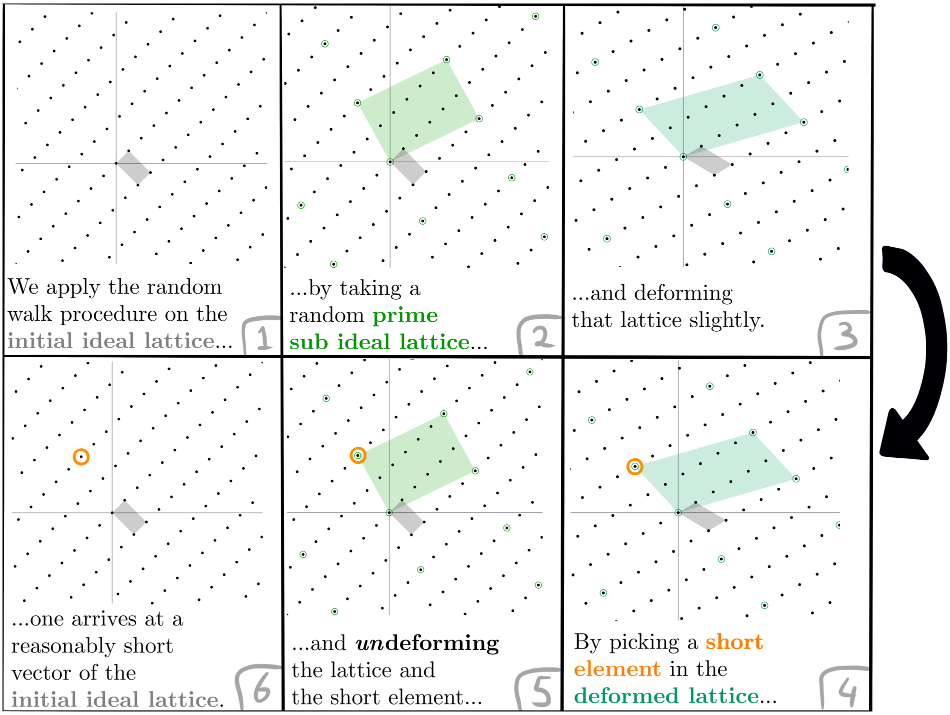

We apply the random walk procedure on the **initial ideal lattice**...

...by taking a random **prime sub ideal lattice**...

...and deforming that lattice slightly.

By picking a **short element** in the **deformed lattice**...

...and ***un*deforming** the lattice and the short element...

...one arrives at a reasonably short vector of the **initial ideal lattice**.

Figure 1.23.: This infographic (note the unusual order of the panels) explains why finding a short element in the lattice at the *end* of a random walk allows to find a reasonably short element in the initial lattice as well. However, there is some loss of shortness quality: the orange element is the *shortest* (non-zero) element in the deformed lattice, but it is only a *reasonably short* element in the **initial lattice**. Summarizing, the random walk does indeed relate the shortest vector problem in two different lattices, but with a slight loss of shortness quality, which is about $\sqrt{n}$ in degree n cyclotomic fields.

subtlety in this picture. Though, because the difference in parameters is rather small in most fields[5], we chose to phrase the simplified statement as a comparison of the same shortest vector problem on the Arakelov class group.

1.7. Ideal Sampling

Introduction

Another application of the random walk on Arakelov class groups allows for efficient sampling of (almost) *prime ideals*. This efficient sampling can be used to compute power residue symbols in polynomial time, assuming the Riemann hypothesis for Hecke-L functions.

Density of prime ideals

The prime number theorem states that the number of primes below a given bound X is roughly equal to $X/\log(X)$. Something similar is true for prime ideals in number fields: the number of prime ideals with norm below X is also roughly equal to $X/\log(X)$, a fact known as *Landau's prime ideal theorem*. Formally,

$$|\{\mathfrak{p} \text{ prime ideal of } \mathcal{O}_K \mid \mathcal{N}(\mathfrak{p}) \leq X\}| \approx X/\log(X). \tag{1.2}$$

Note that this estimated number $X/\log(X)$ of prime ideals with bounded norm does *not* depend on the number field. It seems that all number fields have about the same number of prime ideals with norm below some given X. However, the number of *all* integral ideals with bounded norm *does* vary

[5] The loss in shortness quality in generic number fields K is $O(n \cdot |\Delta_K|^{\frac{1}{2n}})$, where Δ_K is the discriminant of the field. For number fields relevant for cryptography (which have discriminants that grow at most exponential in the degree) this is polynomially bounded in the degree n.

among different number fields. We namely have the following asymptotic estimate:

$$|\{\mathfrak{a} \text{ integral ideal of } \mathcal{O}_K \mid \mathcal{N}(\mathfrak{a}) \leq X\}| \approx \rho_K \cdot X. \tag{1.3}$$

In other words, the number of integral ideals with norm bounded by X grows linearly in X, with slope $\rho_K = \lim_{s\to 1}(s-1)\zeta_K(s)$, the residue at $s = 1$ of the Dedekind zeta function of the concerned field.

By dividing Equation (1.2) by Equation (1.3), one obtains the average number of prime ideals among all ideals. This quantity can be considered as the *density of prime ideals* among all integral ideals, which then roughly equals

$$1/(\rho_K \cdot \log(X)).$$

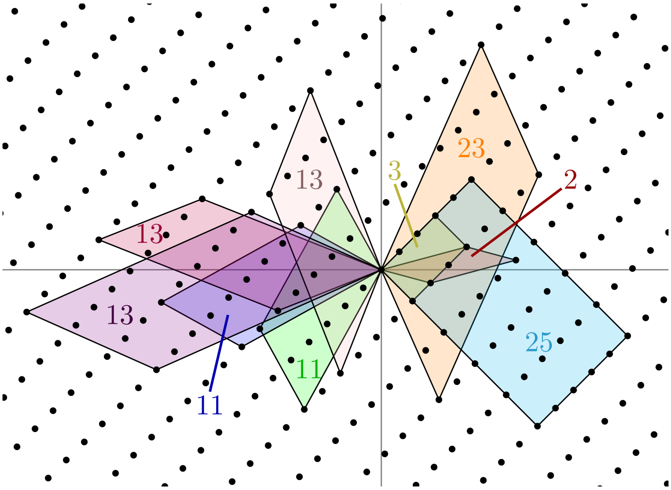

Figure 1.24.: In this image, all shapes of the prime ideal lattices of the number field $\mathbb{Q}(\sqrt{3})$ with norm (i.e., surface area) below 25 are portrayed, with their respective surface area. There are nine such prime ideal lattices. One can see that 2 and 3 ramify, 11, 13 and 23 totally split and 5 is inert in this number field.

Sampling primes

Intuitively, this density estimate gives an algorithm idea to obtain prime ideals in number fields. Namely, sample a *random ideal* with norm below X, and check whether it is prime or not. By this density estimate, the success probability is about $1/(\rho_K \cdot \log(X))$, which is inverse polynomial in the size of X, if we ignore ρ_K for the moment.

In this thesis, we give an *ideal sampling algorithm* that precisely allows this sampling of random ideals, in such a way that indeed the probability of sampling a prime ideal equals $1/(\rho_K \cdot \log(X))$. This technique involves a uniformly random distribution on the Arakelov group.

> Let $\mathfrak{a}$ be an ideal whose Arakelov class is uniformly random distributed, and let $\alpha \in \mathfrak{a} \cap [-r, r]^n$ be uniformly sampled from those elements in $\mathfrak{a}$ that lie in the box $[-r, r]^n$.
>
> Then the probability that the ideal $(\alpha) \cdot \mathfrak{a}^{-1}$ is a prime ideal is at least $1/(3 \cdot \rho_K \cdot \log(r^n))$.

In this statement, there is a necessity for $\mathfrak{a}$ to be randomly distributed on the Arakelov class group, which is absolutely not the case for any fixed ideal $\mathfrak{b}$. But by means of the random walk procedure on the Arakelov class group, one can make *any* fixed ideal $\mathfrak{b}$ 'random' by multiplying it by sufficiently many random small prime ideals and apply a slight deformation, yielding $\mathfrak{a} = x \cdot \prod_j \mathfrak{p}_j \cdot \mathfrak{b}$. This ideal is very close to randomly distributed on the Arakelov class group.

In this way we can algorithmically make $\mathfrak{b}$ randomly distributed, but something is lost as well. Omitting the deformation for the sake of simplicity, sampling $\alpha \in \mathfrak{a} = \prod_j \mathfrak{p}_j \cdot \mathfrak{b}$ gives a guarantee for $(\alpha) \cdot \mathfrak{a}^{-1}$ to be prime with a certain probability. But the fraction $(\alpha) \cdot \mathfrak{b}^{-1}$ can only be guaranteed to be a prime 'up to' the small primes $\prod_j \mathfrak{p}_j$. For most applications, though, this does not cause serious obstacles.

Applications

One of the applications of this prime sampling procedure is that it allows to compute *power residue symbols* in cyclotomic fields $\mathbb{Q}(\zeta_m)$.

The power residue symbol is a function $(\frac{\alpha}{\mathfrak{b}})$ with input $\alpha \in \mathbb{Q}(\zeta_m)$ and $\mathfrak{b} \subseteq \mathbb{Z}[\zeta_m]$ that outputs a power ζ_m^j of the m-th root of unity. It satisfies the properties

(i) $\left(\frac{\alpha}{\beta}\right) = 1$ for $\beta \equiv 1$ modulo $m^m\alpha$;
(ii) $(\frac{\alpha}{\mathfrak{b}\mathfrak{c}}) = (\frac{\alpha}{\mathfrak{b}}) \cdot (\frac{\alpha}{\mathfrak{c}})$, that is, multiplicativity in the lower input;
(iii) $\left(\frac{\alpha}{\mathfrak{p}}\right)$ (with a prime ideal in the lower input) is efficiently computable.

One can make use of these three properties in the following way. To compute $(\frac{\alpha}{\mathfrak{b}})$, apply a random walk on $\mathfrak{b}$, yielding $\tilde{\mathfrak{b}} = \prod_j \mathfrak{p}_j \mathfrak{b}$ and sample $\beta \in \tilde{\mathfrak{b}}$ (omitting the deformation for simplicity). Then $\beta \cdot \tilde{\mathfrak{b}}^{-1} = \mathfrak{p}$ is a prime with good probability. Slightly modifying the sampling procedure, one can assume that β satisfies $\beta \equiv 1$ modulo $m^m\alpha$. By subsequently using properties (i), (ii) and (iii) of the power residue symbol, one obtains an efficiently computable expression for $(\frac{\alpha}{\mathfrak{b}})$.

$$1 = \left(\frac{\alpha}{\beta}\right) = \left(\frac{\alpha}{\mathfrak{p}\prod\tilde{\mathfrak{b}}}\right) = \underbrace{\left(\frac{\alpha}{\mathfrak{p}}\right) \cdot \prod_j \left(\frac{\alpha}{\mathfrak{p}_j}\right)}_{\substack{\text{efficiently computable}\\ \text{by property (iii)}}} \cdot \left(\frac{\alpha}{\mathfrak{b}}\right),$$

The modification of the sampling procedure in order to have $\beta \equiv 1$ modulo $m^m\alpha$ is not entirely trivial and requires a generalization of the random walk theorem over Arakelov *ray* class groups.

Sampling in other ideal sets

Though in this introduction only the set of prime ideals is considered, any subset of the set of ideals of a number field can be taken in place, accounting for the density of this specific set of ideals. For example, the set of *smooth* ideals, ideals that only have prime divisors with small norm, is also an interesting case, as they play a role in class group and unit group computations.

1.8. The Continuous Hidden Subgroup Problem

One particular subject in this thesis is quite separate from the others: the *continuous hidden subgroup problem*. Though this computational problem does concern (general) lattices, it does not have a very direct relation to Arakelov class groups. The analysis of the continuous hidden subgroup in this thesis is a refinement of that of Eistenträger et al. [Eis+14].

Period-finding

The continuous hidden subgroup problem is about recognizing *periodicity* in a continuous signal. Such a continuous signal can be thought of as a sound signal traveling through the air, and its periodicity is then the frequency or pitch of this sound.

A computer solving this hidden subgroup problem, in this analogy, then resembles a violinist with the ability of absolute pitch: given a sound signal, this violinist directly recognizes it and utters 'B-flat', which is around 233 Hertz.

In reality, a sound signal, especially one from a rich-sounding instrument like a violin, consists not just of one single sine tone. It has a certain timbre, which is characterized by the *harmonics* of the tone. Those harmonics are tones that are simultaneously heard and that have frequencies that are exactly integer multiples of the 'main tone'. In the case of the B-flat of 233 Hertz, for example, the harmonic tones have frequencies $233 \cdot 2 = 466$ Hertz, $233 \cdot 3 = 699$ Hertz, *ad infinitum*, see Figure 1.25.

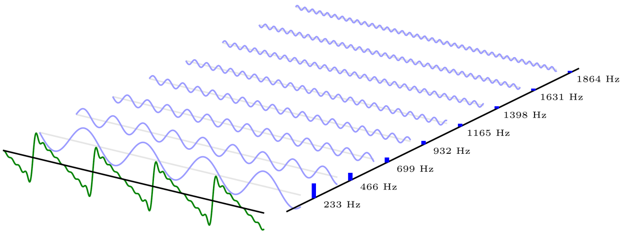

Figure 1.25.: A violin tone has *harmonics*, simultaneously heard tones whose frequency is an integer multiple of the main frequency (233 Hertz, in this example). The variety in loudness of these harmonics defines the timbre.

Period-finding in higher dimensions

A sound signal can be considered one-dimensional, where the one dimension comes from time. Though, the more complex periodicity arises in higher dimensions, since periodicity is then encoded by a lattice, see Figure 1.26.

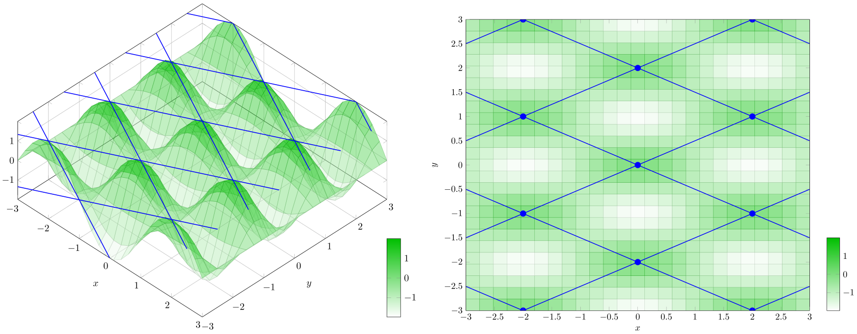

Figure 1.26.: An example of a two-dimensional periodic signal: on the left a 3d-view and on the right a top view. The periodicity can be described by a lattice. The task of the hidden subgroup problem is to retrieve this lattice from the two-dimensional periodic signal.

The higher the dimension of the signal (for our purposes[6], the dimension

[6] One application of the solution of the hidden subgroup problem is in number theory. It can be used to compute unit groups and class groups of number fields [Eis+14]. Also it

does not stop at three), the higher the dimension of the associated periodicity lattice. The 'harmonics' of such multidimensional periodic signal must then be seen as the points of the associated period lattice.

The Fourier transform

The procedure that extracts this periodicity from a signal, including its 'harmonics' (the lattice points), and thus solves the continuous hidden subgroup problem, is called the *Fourier transform*, see Figure 1.27.

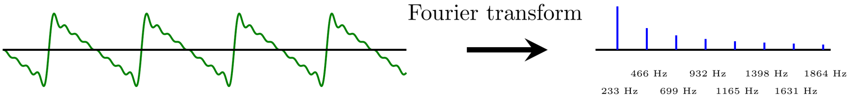

Figure 1.27.: The Fourier transform allows to find the frequencies occurring in a signal, as well as their respective loudness or amplitude.

Though, computers cannot reasonably process a continuous signal as a whole; instead, a computer can only take a finite number of points from the signal. This process is called *discretization*. Due to this discretization, there is some *loss of information* from the signal; the values 'in between' are not known anymore. This particular loss causes the computed Fourier transform of the (discretized) signal to have errors, see Figure 1.28.

Summarizing, by the fact that computers are unable to process infinite continuous signals as a whole, intrinsic errors or 'noise' occurs. If this noise is too large, the out of the computation is unusable.

Errors in the Fourier transform

Whenever the signal is in one dimension, these errors are not that severe. In higher dimensions, though, these errors get *exponentially worse*. This can be

has applications in cryptography, as this solution to the hidden subgroup can also be used to find reasonably short vectors in ideal lattices [Cra+16].

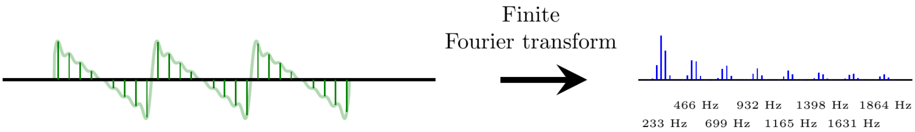

Figure 1.28.: Due to taking only finitely many samples of the periodic signal, small errors occur in the output of the (finite) Fourier transform. In this particular example, the output still resembles the actual frequencies of the original signal (see Figure 1.27), but if there were less sampling points, the output would be so noisy that it would be unusable.

considered as an example of the *curse of dimensionality*, a general expression for describing computational difficulties whenever spatial dimensions grow.

As a consequence, to counteract the explosion of the error size, the number of *samples* of the signal need to grow exponentially as well. This causes this solution for the continuous hidden subgroup problem using Fourier transforms not to be feasible for a normal, classical computer. Instead, we need to use a *quantum* computer.

The Quantum Fourier transform

Due to the special recursive nature of the Fourier transform, it can be efficiently computed by a quantum computer, even when an exponential number of samples is required[7]. In this thesis, in Chapter 3, a thorough analysis is made of how many quantum resources are needed in order to keep the exponentially growing error manageable, depending on properties of the high-dimensional periodic signal. For example, the number of qubits

[7] In reality, these samples are queried in parallel, by using *quantum parallelism*, which allows to sample an exponential number of samples in a parallel way, using only a polynomial amount of classical and quantum resources (i.e., qubits and quantum gates). Also, the output of a quantum Fourier transform yields a quantum state whose *amplitudes* contain the values of the Fourier transform, whose are thus inaccessible due to the nature of quantum phenomena. Fortunately, in this particular hidden subgroup problem, we are only interested in the frequencies where those amplitudes are *high*; such frequencies can then be obtained by *measuring* the quantum state.

(quantum bits) depends logarithmically on how rapidly the signal oscillates and how small one would like the error caused by the discretization to be.

The continuous hidden subgroup problem in higher dimensions, which consists of finding the *hidden period lattice* of a periodic high-dimensional signal, can be solved efficiently on a quantum computer. For an appropriate choice of quantum resources, the errors induced by discretization (i.e., taking only finitely many samples of the signal) can be shown to be feasibly small.

1.9. Outline and Contributions of this Thesis

After this introductory chapter, this dissertation proceeds with Chapter 2, the preliminaries: it states and concisely covers knowledge that is expected from the reader before continuing with the actual results of this thesis.

The next chapter, Chapter 3, is about the continuous hidden subgroup problem, and more or less stands on its own. The contributions of this chapter have been published in the following article, in a slightly different form.

Koen de Boer, Léo Ducas, Serge Fehr. On the Quantum Complexity of the Continuous Hidden Subgroup Problem. In *Advances in Cryptology – EUROCRYPT 2020* [BDF20].

The subsequent chapter, Chapter 4 is about random walks on the Arakelov ray class group. The contributions of this chapter have been published in Section 3 of the following paper, though only for Arakelov class groups with a trivial modulus $\mathfrak{m} = \mathcal{O}_K$. The generalization to arbitrary moduli in this dissertation is new.

Koen de Boer, Léo Ducas, Alice Pellet-Mary, Benjamin Wesolowski. Random Self-reducibility of Ideal-SVP via Arakelov Random Walks. In *Advances in Cryptology – CRYPTO 2020* [Boe+20].

Chapter 5 is about an application of the random walk theorem: a worst-case to average-case reduction for Hermite-SVP on ideal lattices. The contributions in this chapter have been published as well in the CRYPTO 2020 [Boe+20] paper above, with minor differences in some of the proofs concerning discretization.

The last two chapters, Chapter 6 about ideal sampling and Chapter 7 about provably computing the power residue symbol, contain results that have not been published yet.

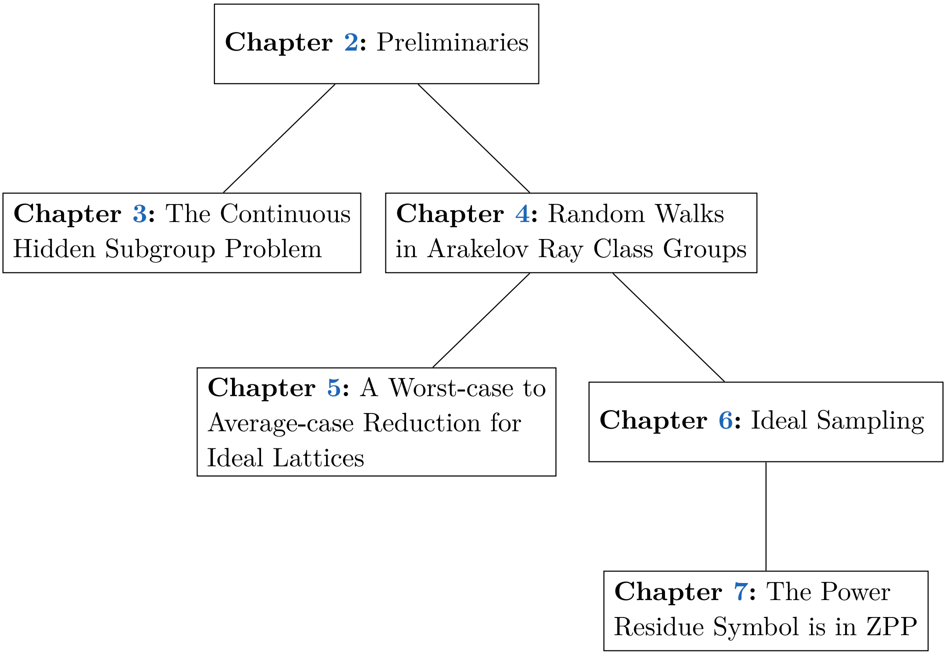
Chapter 2: Preliminaries
Chapter 3: The Continuous Hidden Subgroup Problem
Chapter 4: Random Walks in Arakelov Ray Class Groups
Chapter 5: A Worst-case to Average-case Reduction for Ideal Lattices
Chapter 6: Ideal Sampling
Chapter 7: The Power Residue Symbol is in ZPP

Figure 1.29.: In this diagram is depicted how the chapters depend on each other content-wise.