



Random walks on Arakelov class groups

Boer, K. de

Citation

Boer, K. de. (2022, September 22). *Random walks on Arakelov class groups*. Retrieved from <https://hdl.handle.net/1887/3463719>

Version: Publisher's Version

[Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

License: <https://hdl.handle.net/1887/3463719>

Note: To cite this publication please use the final published version (if applicable).

Random Walks on Arakelov Class Groups

Proefschrift

ter verkrijging van
de graad van doctor aan de Universiteit Leiden,
op gezag van rector magnificus prof. dr. ir. H. Bijl,
volgens besluit van het college voor promoties
te verdedigen op donderdag 22 september 2022

klokke 11.15 uur

door

Koen de Boer

geboren te Nijmegen, Nederland,
in 1991

Promotores:

Prof. dr. L. Ducas (CWI Amsterdam & Universiteit Leiden)
Prof. dr. R. Cramer (CWI Amsterdam & Universiteit Leiden)

Co-promotor:

Dr. B. Wesolowski (Université de Bordeaux)

Promotiecommissie:

Prof. dr. F.A.J. de Haas (Universiteit Leiden)
Dr. E. Kirshanova (Technology Innovation Institute, Abu Dhabi & Immanuel Kant Baltic Federal University, Kaliningrad)
Prof. dr. R.M. van Luijk (Universiteit Leiden)
Prof. dr. R. Schoof (Università di Roma “Tor Vergata”)
Prof. dr. D. Stehlé (École Normale Supérieure de Lyon)

The research was carried out in the Cryptology Group at CWI Amsterdam. The PhD position was funded by the ERC Advanced Grant 740972 (ALGSTRONGCRYPTO) and by the European Union Horizon 2020 Research and Innovation Program Grant 780701 (PROMETHEUS).



Centrum Wiskunde & Informatica



Universiteit
Leiden

Koen de Boer

Random Walks on Arakelov Class Groups
and Their Applications in Cryptography
and Algorithmic Number Theory

To my family.

Contents

1. Introduction	1
1.1. Number Theory	1
1.2. Ideal Lattices	4
1.3. Arakelov Class Groups	7
1.4. Random Walks on the Arakelov Class Group	13
1.5. The Random Walk Theorem for Arakelov Class Groups	19
1.6. A Worst-case to Average-case Reduction	25
1.7. Ideal Sampling	30
1.8. The Continuous Hidden Subgroup Problem	34
1.9. Outline and Contributions of this Thesis	38
2. Preliminaries	41
2.1. General Notation	41
2.2. Fourier Theory	41
2.2.1. Groups	42
2.2.2. Norms and Fourier Transforms	43
2.2.3. The Poisson Summation Formula	45
2.2.4. The Fourier Transform of Vector-valued Functions	48
2.2.5. Trigonometric Approximation	51
2.3. Number Theory	53
2.3.1. Algebraic Number Theory	53
2.3.2. The Extended Riemann Hypothesis	56
2.3.3. Prime Densities	57
2.4. Arakelov Theory	59
2.4.1. The Arakelov Ray Divisor Group	59
2.4.2. The Arakelov Ray Class Group	60
2.4.3. Relation with Other Number-theoretic Groups	62

2.4.4. The Volume of the Arakelov Ray Class Group	64
2.4.5. An Example of an Arakelov Class Group	66
2.5. Lattices	72
2.5.1. General Lattices	72
2.5.2. Divisors and Ideal Lattices	72
2.5.3. The Gaussian Function and Smoothing Errors	76
2.5.4. Gaussian Distributions	79
2.6. The Lipschitz Condition	79
3. The Continuous Hidden Subgroup Problem	81
3.1. Summary	81
3.2. Introduction	84
3.3. Problem Statements and Results	90
3.3.1. Notation and Set-up	90
3.3.2. Main Theorem: Continuous Hidden Subgroup Problem	91
3.3.3. Dual Lattice Sampling Problem	93
3.3.4. Full Dual Lattice Recovery	94
3.3.5. Primal Basis Reconstruction	94
3.3.6. Gaussian State Preparation	95
3.3.7. Proof of the Main Theorem	96
3.4. Dual Lattice Sampling Algorithm	99
3.4.1. The Algorithm	99
3.4.2. The Figure of Merit	102
3.5. Analysis	104
3.5.1. Proof Overview	104
3.5.2. Formal Analysis	107
3.6. From Sampling to Full Dual Lattice Recovery	117
3.7. Recovering a Basis of the Primal Lattice	123
3.7.1. An Approximate Well-conditioned Basis of the Dual	124
3.7.2. Inverting the Dual Approximate Basis	126
3.7.3. Combining the Errors and Tuning the Parameters .	127
4. Random Walks on Arakelov Ray Class Groups	129
4.1. Summary	129
4.2. Introduction	131

4.3. Random Walk Theorem for the Arakelov Ray Class Group	138
4.3.1. Main result	140
4.3.2. Hecke Operators	142
4.3.3. Bounds on Eigenvalues of Hecke Operators	143
4.3.4. The Infinite Analytic Conductor	147
4.3.5. Fourier Analysis on the Ray Unit Torus	153
4.3.6. Splitting up the Character Decomposition	154
4.3.7. Conclusion	158
5. A Worst-case to Average-case Reduction for Ideal Lattices	165
5.1. Summary	165
5.2. Introduction	166
5.2.1. The Result	167
5.2.2. Overview	168
5.2.3. Related work	169
5.3. Representation of Ideal Lattices by Means of Distributions	171
5.4. The Worst-case to Average-case Reduction	179
5.5. Discretizing the Reduction Algorithm	189
5.5.1. Introduction	189
5.5.2. Precise Definition of the Distributions Involved	191
5.5.3. Discretized Algorithm Analogues	194
5.5.4. Closeness Proofs	198
6. Ideal sampling	203
6.1. Summary	203
6.2. Introduction	204
6.2.1. Our Technique	204
6.2.2. Applications	206
6.2.3. Related Works	207
6.3. Preliminaries	207
6.4. Probability-density Correspondence	209
6.4.1. Result	209
6.4.2. Proof Overview of Theorem 6.9	211

6.5. Extended Proof of Theorem 6.9	216
6.5.1. Simplify the Probability by Fixing a Single Ideal $\mathfrak{c} \in \mathcal{S}^m$ and a single Arakelov divisor $\mathbf{a} \in \text{Div}_{K^m}^0$	216
6.5.2. Estimating the Number of Shifted Lattice Points in a Box	217
6.5.3. Estimating the Probability of Sampling a Single Fixed Ideal for a <i>Random</i> Arakelov Divisor	220
6.5.4. The Number $ \text{Exp}(\mathbf{a} + d(\mathfrak{c}))_\tau^\times \cap r\mathcal{B}_\infty $ Only Depends on the Arakelov Ray <i>Class</i> of $\mathbf{a} \in \text{Div}_{K^m}^0$	220
6.5.5. Taking the Logarithmic Map into $H = \text{Log } K_{\mathbb{R}}^0$	223
6.5.6. Applying the Abel Summation Formula	224
6.6. Ideal Sampling	226
6.6.1. Sampling in a Box	226
6.6.2. The Sampling Algorithm	228
7. The Power Residue Symbol is in ZPP	235
7.1. Summary	235
7.2. Introduction	236
7.3. Preliminaries	238
7.4. Reduction to Cyclotomic Fields	241
7.4.1. Introduction	241
7.4.2. Proof Strategy	242
7.4.3. Signature Identity	244
7.4.4. Invariant Factor Decomposition of R/\mathfrak{b}	248
7.4.5. The Signature is Compatible with Short Exact Sequences	249
7.4.6. The Determinant Formula	251
7.4.7. Applying Induction on the Components of the Module	253
7.4.8. Conclusion	254
7.5. Computing the Power Residue Symbol in Cyclotomic Fields	255
7.5.1. Main Idea	255
7.5.2. The Full Algorithm	257

7.6. Discussion	259
7.6.1. Computing Hilbert Symbols Using Power Residue Symbols	259
7.6.2. Computing Artin Symbols in the Same Fashion	261
Bibliography	265
A. Appendix	279
A.1. Number-theoretic Computations	279
A.2. Bound on the Residue of the Zeta Function for Cyclotomic Fields	283
A.2.1. Splitting $\log(\rho_K) = R_K + M_K$ into a Ramified Term and a Main Term	285
A.2.2. Estimating the Ramified Term	286
A.2.3. Splitting the Main Term in an Initial Part and a Tail Part	288
A.2.4. Estimating the Initial Part of the Main Term	289
A.2.5. Estimating the Tail Part of the Main Term	291
A.3. Exact Sequences	292
A.4. The Yudin-Jackson Theorem	294
A.5. The Gaussian State	299
A.5.1. Reducing to the One-dimensional Case	299
A.5.2. The Periodic and Non-periodic Discrete Gaussian	300
A.5.3. Computing the Periodic Gaussian State	301
A.5.4. Estimating the Complexity and Fidelity	303
A.5.5. Proof of Lemma A.27	306
A.6. Discrete Gaussians	308
Summary	313
Samenvatting	315
Acknowledgments	317
Curriculum Vitae	319

Contents

List of Symbols 321

Index 329

List of Figures

1.1. The number ring $\mathbb{Z}[\sqrt{2}]$ visualized on the real plane	3
1.2. The Minkowski embedding of $\mathbb{Z}[\sqrt[3]{2}]$ into $\mathbb{R} \times \mathbb{C}$	4
1.3. Examples of ideal lattices and non-ideal lattices	5
1.4. Deformation of ideal lattices	6
1.5. Examples of trivial class and non-trivial class ideal lattices .	7
1.6. Ideal lattice shapes shifting seamlessly from  to 	8
1.7. Examples of ideal lattices in $\mathbb{Q}(\sqrt{3})$	9
1.8. Periodicity of the ideal lattices in $\mathbb{Q}(\sqrt{3})$	10
1.9. Circularity of the ideal lattices in $\mathbb{Q}(\sqrt{3})$	11
1.10. The Arakelov class group of a number field K consists of a union of finitely many hypertori.	12
1.11. Ideal lattices on the same torus can be transformed into each other	13
1.12. Random walk analogy with an ant	14
1.13. Gaussian distribution ‘folding around’ the torus	15
1.14. Random walks on Arakelov class groups require ‘jumps’ .	16
1.15. Examples of prime sub ideal lattices	18
1.16. Examples of random jumps on the Arakelov class group .	18
1.17. Infographic explaining the random walk procedure	20
1.18. An explanation of the random walk’s parameters	20
1.19. The more jumps happen in the random walk, the less crawling is needed	22
1.20. An example of the equidistribution of primes on the Arakelov class group of $\mathbb{Q}(\sqrt{3})$	24
1.21. An explanation of the Shortest Vector Problem	26

1.22. A question about the hardness of SVP being the same on all ideal lattices	27
1.23. Infographic explaining the worst-case to average-case reduction	29
1.24. Picture of all prime ideal lattices of $\mathbb{Q}(\sqrt{3})$ with norm below 25	31
1.25. Harmonics of a violin tone	35
1.26. An example of a two-dimensional periodic signal and its period lattice	35
1.27. The Fourier transform captures periodicity	36
1.28. Discretization in the Fourier transform causes intrinsic noise	37
1.29. The dependencies of the chapters in this thesis	39
2.1. Groups and their Haar measures	43
2.2. An example of restriction of a function	46
2.3. An example of periodization of a function	47
2.4. An example of \mathbb{Z} -periodization of a Gaussian function	48
2.5. Informal illustration of the Poisson summation formula	49
2.6. The Poisson summation formula applied to a Gaussian function	50
2.7. Wide periodic Gaussians are close to uniform	51
2.8. A commutative diagram involving the Arakelov ray class group	62
2.9. A concrete picture of an Arakelov class group (parallelograms)	70
2.10. A concrete picture of an Arakelov class group (tori)	71
3.1. An example of an (r, ϵ) -separating function	91
3.2. The Lipschitz constant of a (r, ϵ) -separating function	92
3.3. Periodicity of continuous signals versus discrete signals	100
3.4. A visual representation of Algorithm 2	101
3.5. A visual explanation of the target set C	103
3.6. An explanation of the indicator-like function $\iota_C(\ell^*)$	111
3.7. An example of an (R, q) -concentrated distribution	117
3.8. Examples of two distributions not being p -evenly distributed	118
3.9. A depiction of the case distinction of Lemma 3.22	120
4.1. The Arakelov class group can be thought of as $ \text{Cl}_K $ copies of the logarithmic unit torus	130

4.2. The discrete walk on the Arakelov class group	135
4.3. The continuous walk on the Arakelov class group	135
4.4. Volumetric covering argument for the Arakelov class group	136
4.5. The effect of the Hecke operator on Arakelov class group distributions	138
4.6. High and low-frequency characters on the hypertorus T^m . .	140
4.7. Hecke operators diminish non-unit characters	141
4.8. The effect of ‘averaging’ on complex exponential functions .	143
4.9. Characters and their associated dual lattice points	154
4.10. Distributions tend to the uniform distribution under repeated action of the Hecke operator	156
4.11. The larger the Gaussian deviation s is, the less influence it has on the running time of the random walk	163
5.1. A short vector in subideals of \mathfrak{a} is also a reasonably short vector in \mathfrak{a} itself	169
5.2. Deformation of ideal lattices	172
5.3. The discrete hyper-circle	192
6.1. Counting lattice points in a box	219
7.1. The multiplicative action of $\langle \zeta \rangle$ on a residue field R/\mathfrak{p}_5 . .	245
7.2. A depiction of the computation of the signature (ϕ, M) . .	246
7.3. The computation of the signature $(\phi_{1+\zeta}, R/\mathfrak{p}_5)$	248
A.1. The kernel-cokernel exact sequence	294

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* $K \rightarrow \bigoplus_{\sigma:K \rightarrow \mathbb{C}} \mathbb{C}$, $\alpha \mapsto (\sigma(\alpha))_\sigma$. In most of the literature, the codomain of this Minkowski embedding is restricted to $K_{\mathbb{R}} = \{x_\sigma \in \bigoplus_{\sigma:K \rightarrow \mathbb{C}} \mathbb{C} \mid \bar{x}_\sigma = x_{\bar{\sigma}}\}$, 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 component-wise 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.

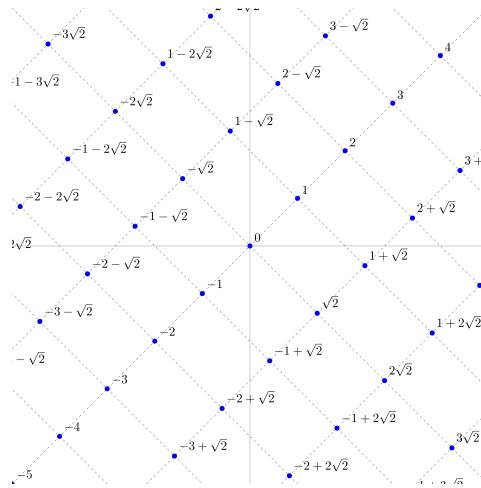


Figure 1.1.: The number ring $\mathbb{Z}[\sqrt{2}]$ visualized on the real plane, using the Minkowski embedding, sending $\sqrt{2} \mapsto (\sqrt{2}, -\sqrt{2})$ and $1 \mapsto (1, 1)$.

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

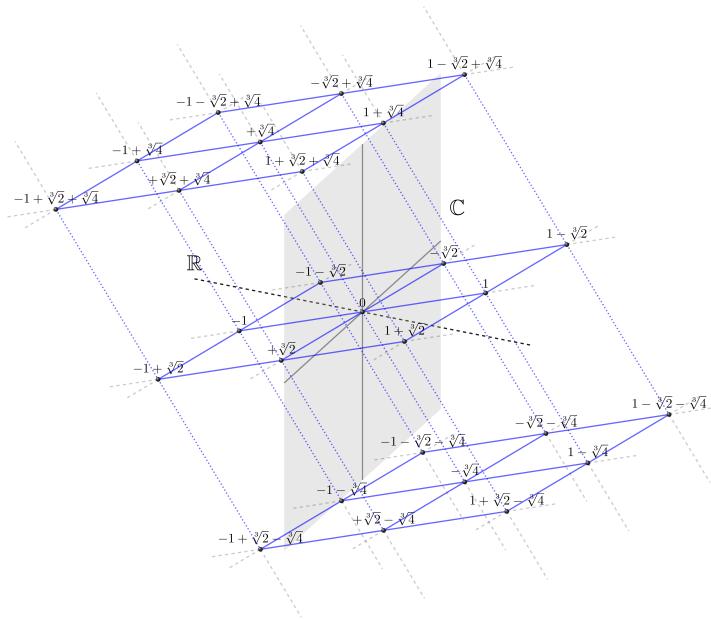


Figure 1.2.: This picture shows the Minkowski embedding of $\mathbb{Z}[\sqrt[3]{2}]$ into $\mathbb{R} \times \mathbb{C}$.

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.

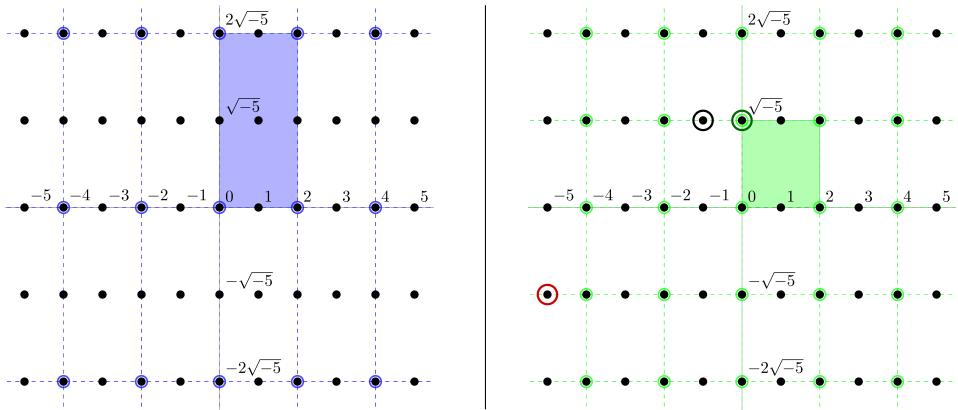


Figure 1.3.: The blue lattice is the ideal lattice $2 \cdot \mathbb{Z}[\sqrt{-5}]$ in $\mathbb{Q}(\sqrt{-5})$ consisting of multiples of 2. The green lattice is an example of a lattice that is *not* an ideal lattice of $\mathbb{Q}(\sqrt{-5})$, because it is not stable under multiplication with elements 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.

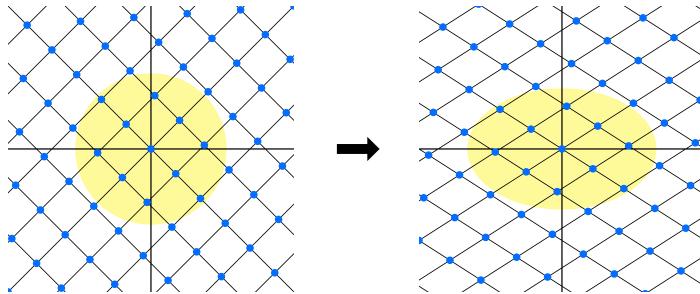


Figure 1.4.: In this two-dimensional example, the left ideal lattice is slightly stretched in the x -direction and slightly squished in the y -direction, leading to the perturbed ideal lattice on the right. The yellow circle functions as a visual aid, making the precise deformation of the lattice more 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 $x \cdot I, y \cdot J \in \text{IdLat}_K^0$ 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.

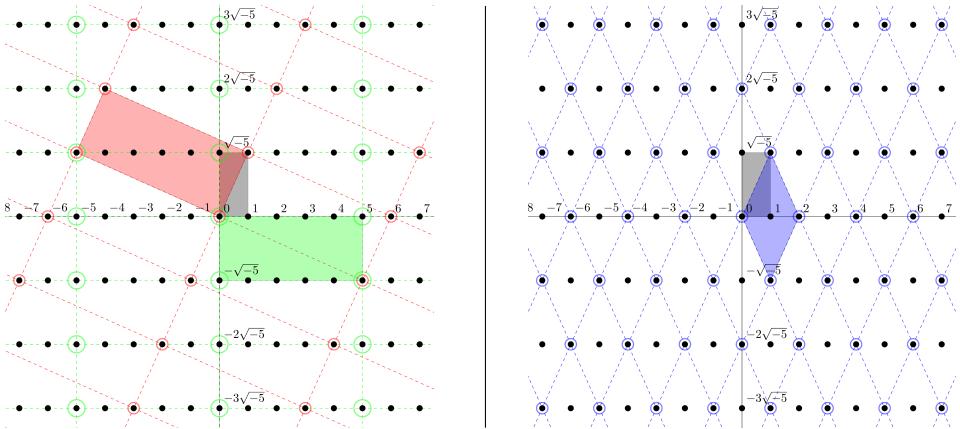


Figure 1.5.: On the left are three trivial-class ideal lattices, namely $\mathbb{Z}[\sqrt{-5}]$, $\sqrt{-5} \cdot \mathbb{Z}[\sqrt{-5}]$ and $(1 + \sqrt{-5}) \cdot \mathbb{Z}[\sqrt{-5}]$. By observing the rectangular shapes enclosed by the lattice points, one indeed observes that these three lattices 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 \blacksquare and the ‘diamond’ class \blacklozenge . This categorization of the ideal lattices

1. Introduction

of $\mathbb{Q}(\sqrt{-5})$ is described by the *Arakelov class group* of $\mathbb{Q}(\sqrt{-5})$, which we denote by $\text{Pic}_{\mathbb{Q}(\sqrt{-5})}^0$. In other words,

$$\text{Pic}_{\mathbb{Q}(\sqrt{-5})}^0 = \{\blacksquare, \blacklozenge\}.$$

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,

$$\text{Pic}_K^0 := \text{IdLat}_K^0 / \sim,$$

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, $|\text{Pic}_{\mathbb{Q}(\sqrt{-5})}^0| = 2$, 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 \blacksquare to \blacklozenge .

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 $\mathbb{Z}[\sqrt{3}]$.

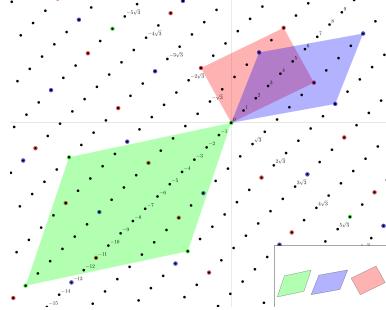


Figure 1.7.: In this picture we examine three different shapes of ideal lattices of the number field $\mathbb{Q}(\sqrt{3})$. The variety of shapes suggests that there is a spectrum of different ideal lattice shapes.

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) embedding¹. 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

¹Technically, 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.

1. Introduction

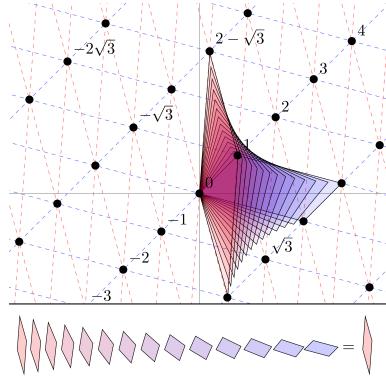


Figure 1.8.: We consider here ideal lattices of the number field $\mathbb{Q}(\sqrt{3})$. By stretching the red shape into the x -direction and shrinking the same amount in the y -direction, one obtains a variety of shapes. Eventually, one arrives at the blue shape, which represents the same ideal lattice 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 $\mathbb{Z}[\sqrt{3}]$.

$t \in \mathbb{R}$,

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

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 $\text{Pic}_K^0 \simeq \mathbb{R}/\log(2 + \sqrt{3}) \cdot \mathbb{Z}$ 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 $H/\text{Log}(\mathcal{O}_K^\times)$ of the hyperplane $H = \text{span}(\text{Log}(\mathcal{O}_K^\times))$ and the logarithmic unit lattice $\text{Log}(\mathcal{O}_K^\times)$* that arises in Dirichlet's unit theorem.

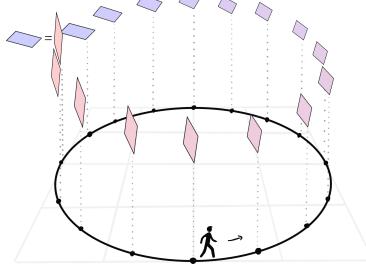


Figure 1.9.: By deforming an initial ideal lattice in $\mathbb{Q}(\sqrt{3})$ appropriately, one eventually arrives at the same ideal lattice. This yields a circular pattern; as a result, the Arakelov class group of $\mathbb{Q}(\sqrt{3})$ is isomorphic to the circle group S^1 .

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 $\text{Cl}(K)$ and the *logarithmic unit torus* $T = H/\text{Log}(\mathcal{O}_K^\times)$, the quotient group of the hyperplane $H = \text{span}(\text{Log}(\mathcal{O}_K^\times))$ and the logarithmic unit lattice $\text{Log}(\mathcal{O}_K^\times)$. Here, $\text{Log}(\eta) := (\log |\sigma(\eta)|)_\sigma$ 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 $\text{Log}(\mathcal{O}_K^\times)$, 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 $\text{Cl}(K)$ and the logarithmic unit torus $T = H/\text{Log}(\mathcal{O}_K^\times)$.

$$0 \rightarrow H/\text{Log}(\mathcal{O}_K^\times) \rightarrow \text{Pic}_K^0 \rightarrow \text{Cl}(K) \rightarrow 0.$$

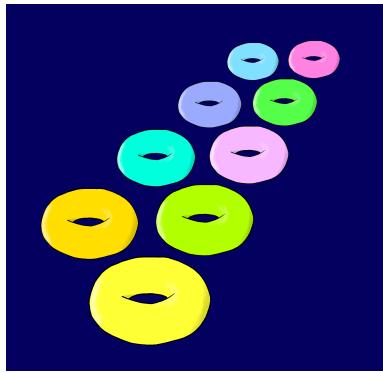


Figure 1.10.: The Arakelov class group of a number field K consists of a union of finitely many hypertori.

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 $T = H/\text{Log}(\mathcal{O}_K^\times)$ 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 $T = H/\text{Log}(\mathcal{O}_K^\times) \simeq \mathbb{R}/\log(2 + \sqrt{3})\mathbb{Z}$, 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 $T = H/\text{Log}(\mathcal{O}_K^\times)$ 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 $T = H/\text{Log}(\mathcal{O}_K^\times)$, 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.

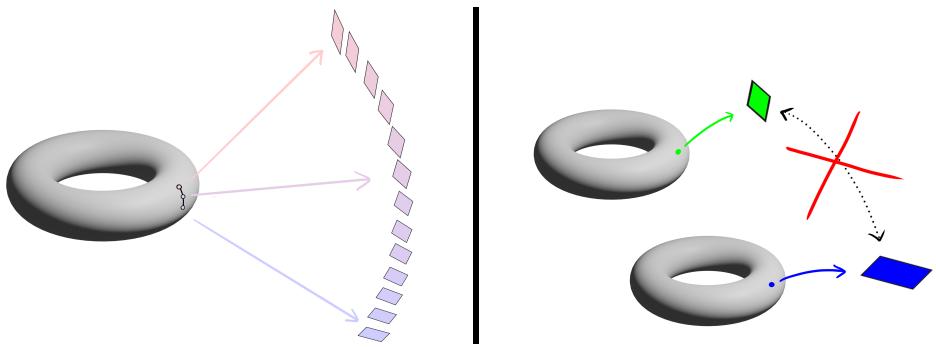


Figure 1.11.: Any two lattices corresponding to points on the *same* torus, can be transformed into each other (left). If two lattices correspond to points on *different* tori, they cannot be transformed into each other (right).

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.

1. Introduction

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).

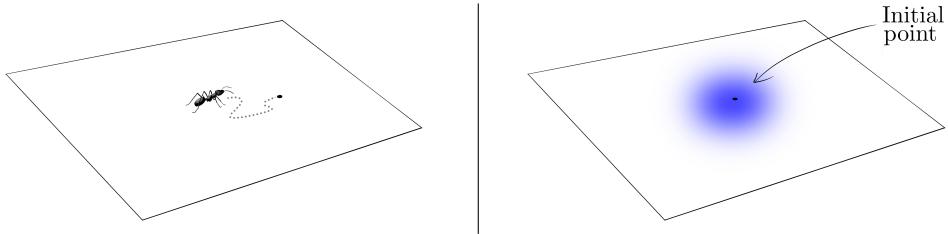


Figure 1.12.: An ant on a plane will, without external stimuli, follow an irregular path, as in the left-most image. This can be regarded as an intuitive interpretation of a random walk. The probability distribution arising from this statistical behavior is called a *random walk distribution* and is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

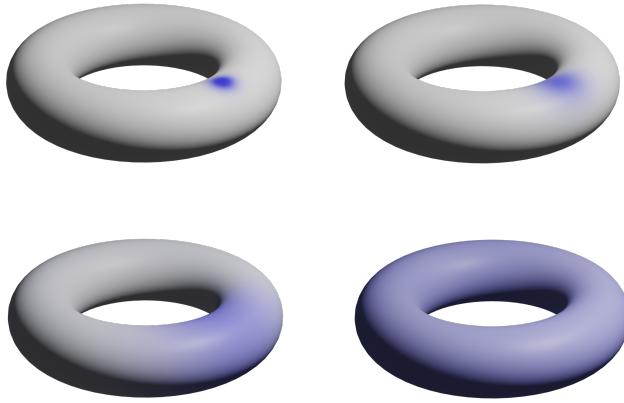


Figure 1.13.: As the deviation of the Gaussian distribution increases, the distribution ‘folds around’ the torus more and more. From a certain value of the deviation, the distribution is very close to a uniform distribution.

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.

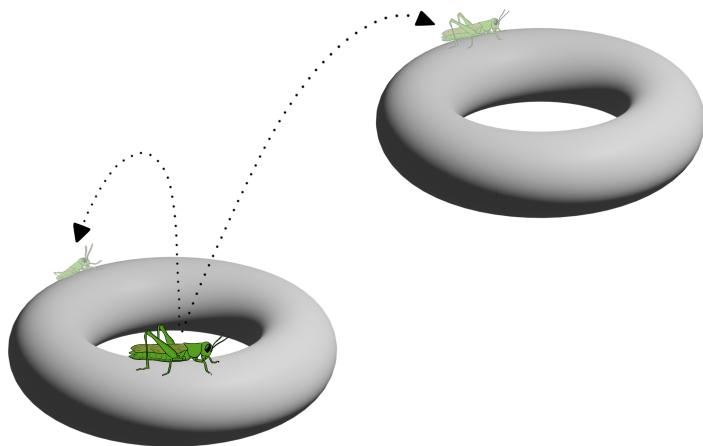


Figure 1.14.: Due to the disconnected nature of the Arakelov class group, as it consists of several separated tori, we also need ‘jumps’ in our random walk, next to ‘crawls’. For intuition it is then more appropriate to have a grasshopper in mind. The grasshopper does not need to land on a different 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 $x \cdot \mathfrak{a} \mapsto x \cdot (\mathfrak{p} \cdot \mathfrak{a})$ yields a jump in the Arakelov class group².

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.

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

1. Introduction

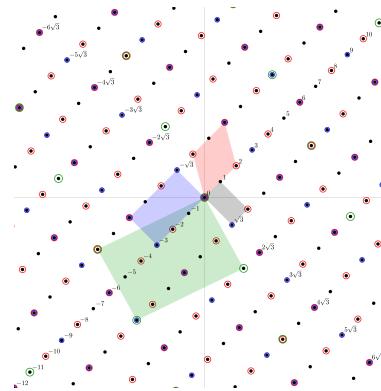


Figure 1.15.: The red, blue and green ideal lattices are all *prime* sub ideal lattices of the gray (base) ideal lattice, because the shapes of the respective ideal lattices are 2, 3 and 11 (prime numbers) times larger than the surface of the gray ideal lattice.

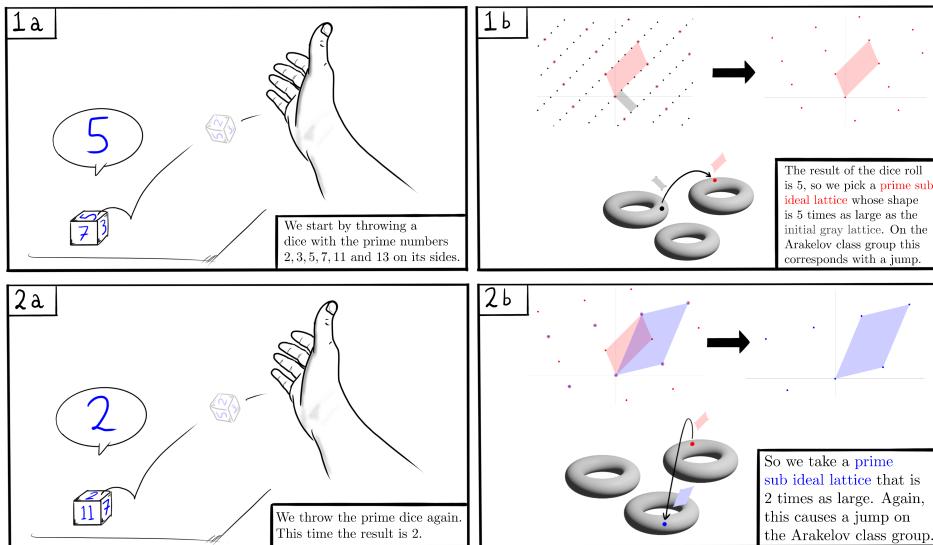


Figure 1.16.: This picture shows two repetitions ($N = 2$) of a random jump. In more realistic cases, both the number of primes and the number of jumps are larger. Note that at each jump, the ideal lattice gets sparser, or, equivalently, its shape gets larger.

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

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:

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

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

1. Introduction

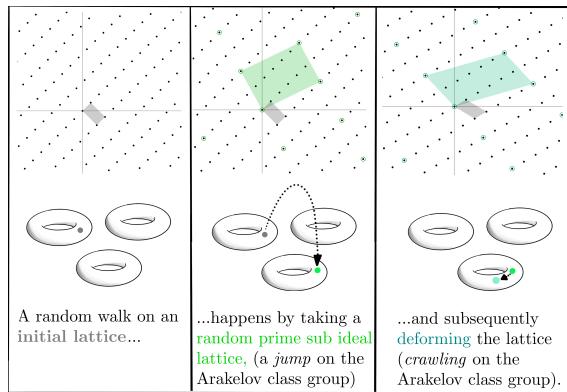


Figure 1.17.: A concrete realization of the random walk procedure on an ideal lattice with a single jump. The prime sub ideal lattice is chosen at random, as well as the deformation (by sampling a Gaussian distribution).

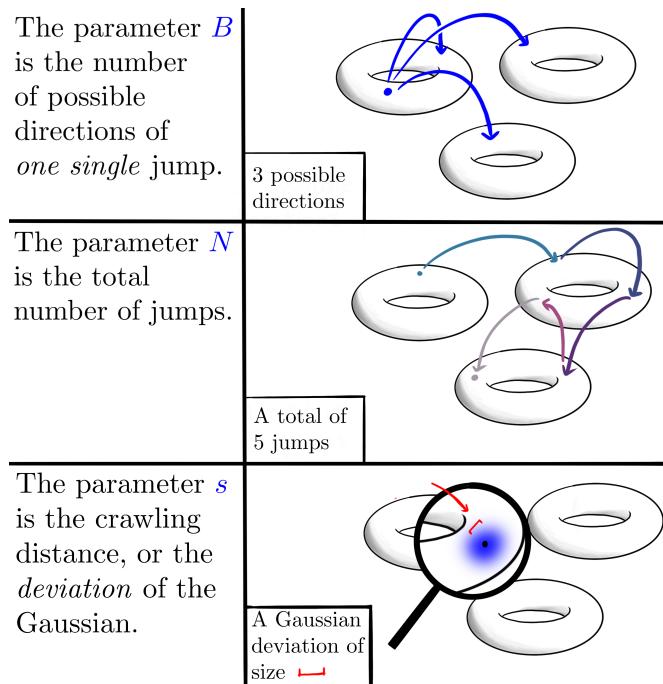


Figure 1.18.: An explanation of the random walk's parameters.

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 primes³ 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 cover⁴ 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 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

³Formally, 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 .

⁴Cover’, 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.

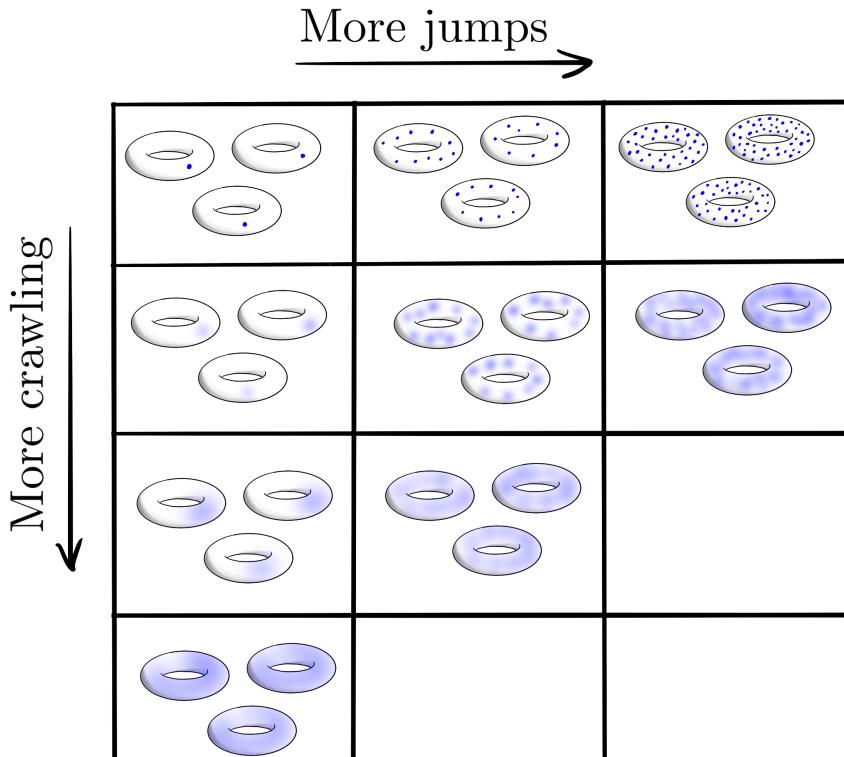


Figure 1.19.: The more jumps happen in the random walk, the less crawling is needed in order to cover the entire Arakelov class group.

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 $\text{vol}(\text{Pic}_K^0)$ of the Arakelov class group, that is,

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

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 $\text{vol}(\text{Pic}_K^0)$. 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

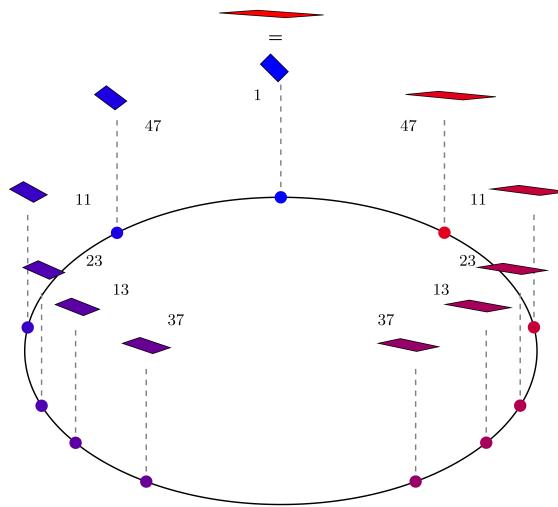


Figure 1.20.: The circle-shaped Arakelov class group of $\mathbb{Q}(\sqrt{3})$ with the positions of the first few prime ideals p and their associated shapes. Already in this ‘small’ example, there is a reasonable equidistribution of these prime ideals on the Arakelov class group.

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 a 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.

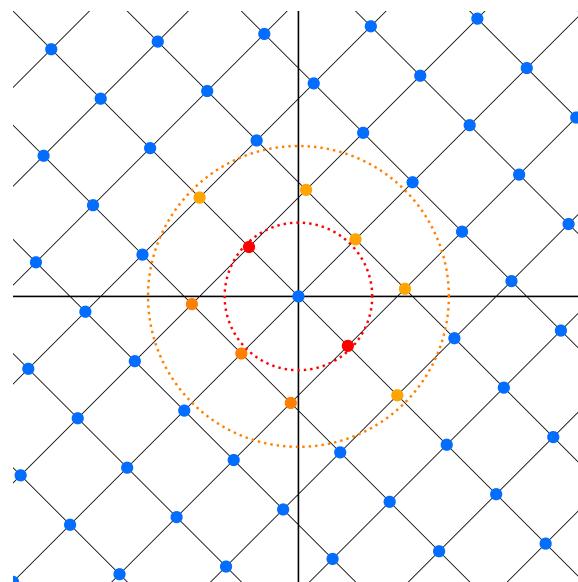


Figure 1.21.: The *shortest vector problem* asks to find a short vector in the lattice, which means that it is close to the origin, but not the origin point itself. The **red points** are the *shortest* lattice elements. In most cases, though, just short vectors, like the **orange points**, are also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

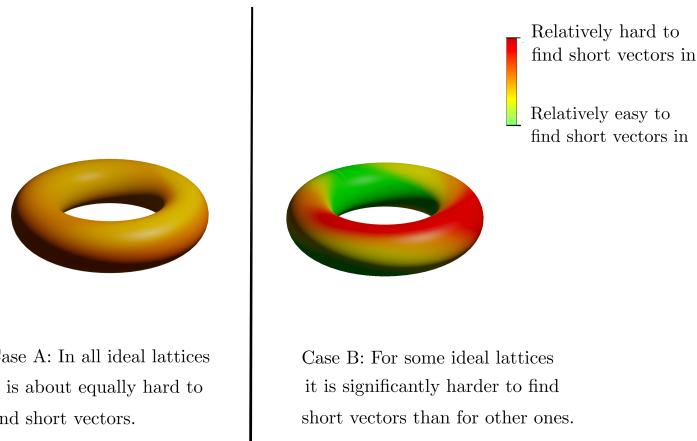


Figure 1.22.: Is it for all ideal lattices on the Arakelov class group about equally hard to find short vectors in (case A) or not (case B)? Note that we just pictured one single torus for the Arakelov class group, for simplicity.

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

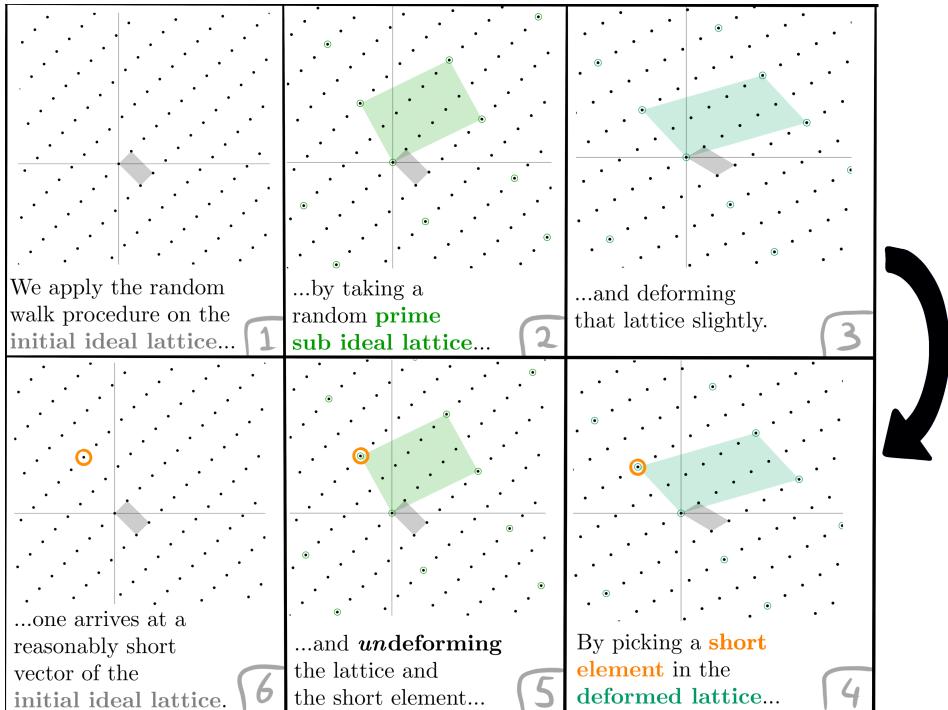


Figure 1.23.: This infographic (note the unusual order of the panels) explains why finding a short element in the lattice at the *end* of a random walk allows to find a reasonably short element in the initial lattice as well. However, there is some loss of shortness quality: the orange element is the *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 fields⁵, 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,

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

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

⁵The 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:

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

In other words, the number of integral ideals with norm bounded by X grows linearly in X , with slope $\rho_K = \lim_{s \rightarrow 1} (s-1) \zeta_K(s)$, 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

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

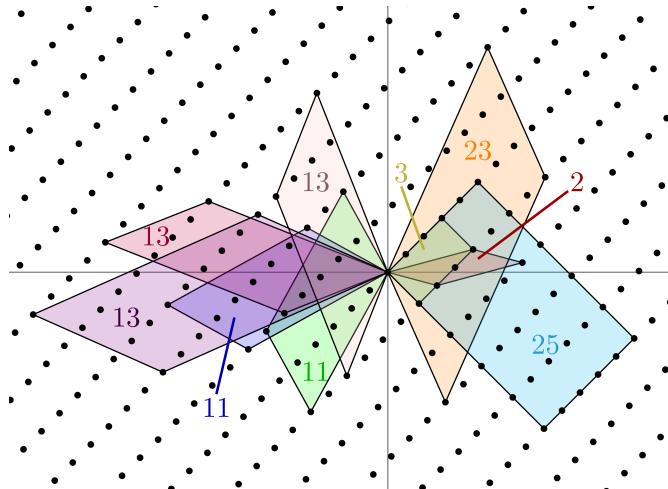


Figure 1.24.: In this image, all shapes of the prime ideal lattices of the number field $\mathbb{Q}(\sqrt{3})$ with norm (i.e., surface area) below 25 are portrayed, with their respective surface area. There are nine such prime ideal lattices. One can see that 2 and 3 ramify, 11, 13 and 23 totally split 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) $(\frac{\alpha}{\beta}) = 1$ 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) $(\frac{\alpha}{\mathfrak{p}})$ (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.

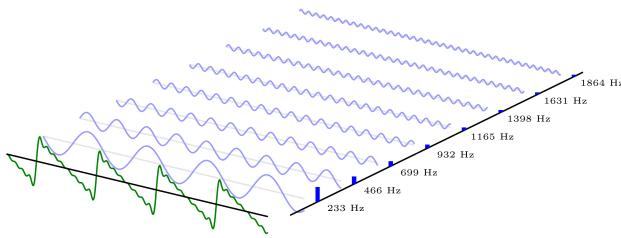


Figure 1.25.: A violin tone has *harmonics*, simultaneously heard tones whose frequency is an integer multiple of the main frequency (233 Hertz, in this example). The variety in loudness of these harmonics defines the timbre.

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.

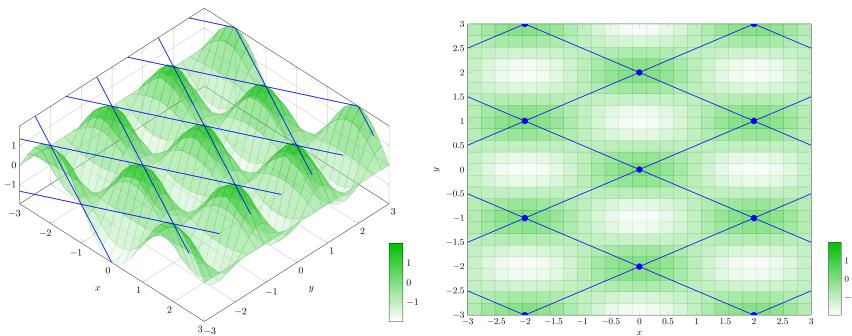


Figure 1.26.: An example of a two-dimensional periodic signal: on the left a 3d-view and on the right a top view. The periodicity can be described by a lattice. The task of the hidden subgroup problem is to retrieve this lattice from the two-dimensional periodic signal.

The higher the dimension of the signal (for our purposes⁶, the dimension

⁶One 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].

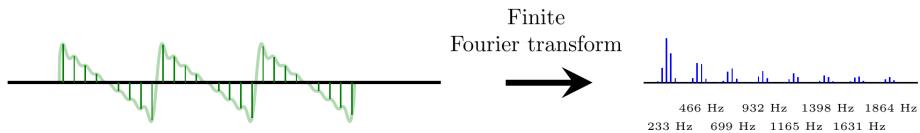


Figure 1.28.: Due to taking only finitely many samples of the periodic signal, small errors occur in the output of the (finite) Fourier transform. In this particular example, the output still resembles the actual frequencies of the original signal (see Figure 1.27), but if there were less sampling points, 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 required⁷. 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

⁷In 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.

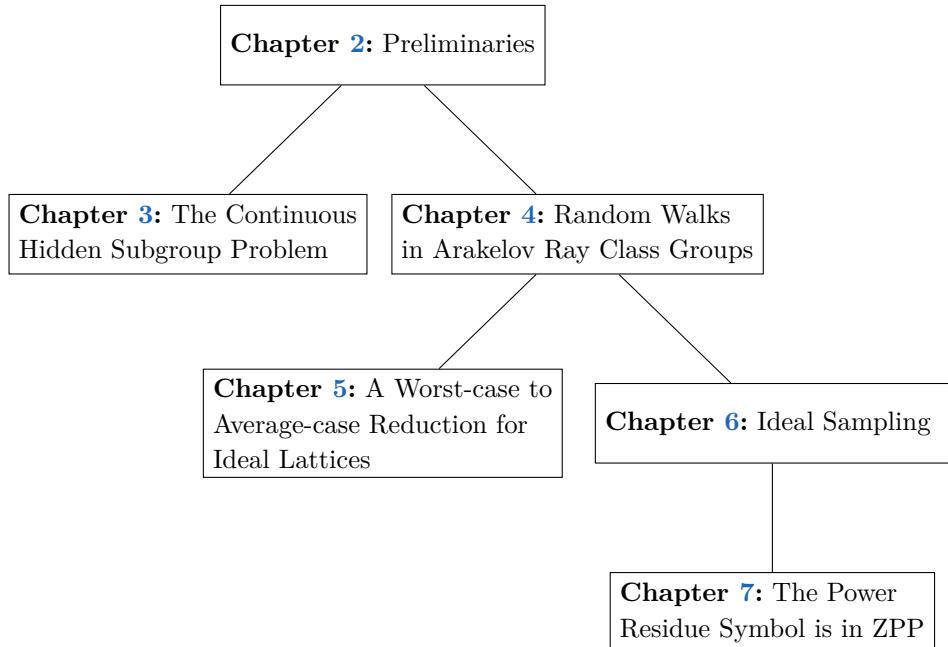


Figure 1.29.: In this diagram is depicted how the chapters depend on each other content-wise.

2. Preliminaries

2.1. General Notation

We denote by $\mathbb{N}, \mathbb{Z}, \mathbb{Q}, \mathbb{R}, \mathbb{C}$ the natural numbers, the integers, the rationals, the real numbers and the complex numbers respectively. All logarithms are in base e . For a rational number $p/q \in \mathbb{Q}$ with p and q coprime, we let $\text{size}(p/q)$ refer to $\log|p| + \log|q|$. We extend this definition to vectors and matrices of rational numbers, by taking the sum of the sizes of all the coefficients.

Vectors $\mathbf{v} \in V$ are denoted in boldface and are to be interpreted column-wise unless stated otherwise. In the case of a vector in a (quantum) Hilbert space \mathcal{H} , we sometimes deviate from this notation and use the bra-ket notation as well; $|\mathbf{v}\rangle$ for primal vectors and $\langle \mathbf{v}|$ for dual vectors. An inner product of $|\mathbf{w}|$ and $|\mathbf{v}\rangle$ is then denoted $\langle \mathbf{w}|\mathbf{v}\rangle$, and the notation for the tensor product $|\mathbf{w}\rangle \otimes |\mathbf{v}\rangle$ of two vectors in a Hilbert space is generally suppressed, i.e., we denote $|\mathbf{w}\rangle|\mathbf{v}\rangle$ instead.

2.2. Fourier Theory

We start with a brief introduction to Fourier analysis over arbitrary locally compact abelian groups. This general treatment allows us to then apply the general principles to the different groups that play a role in this thesis, especially in Chapter 3. For the reader that is unfamiliar with such a general treatment, it is useful—and almost sufficient—to think of \mathbb{R} , of $\mathbb{T} = \mathbb{R}/\mathbb{Z}$,

and a finite group. For more details and for the proofs we refer to Deitmar's book on this subject [DE16].

2.2.1. Groups

Here and below we consider a *locally compact abelian* group G . Such a group admits a *Haar measure* μ that is unique up to a normalization factor. The crucial property of such a Haar measure is that it is invariant under the group action. Simple examples are $G = \mathbb{R}$ with μ the Lebesgue measure λ , or a finite group G with μ the counting measure $\#$.

The *dual group* \hat{G} , consisting of the continuous¹ group homomorphisms χ from G into S^1 , the multiplicative group of complex numbers of absolute value 1, is again a locally compact abelian group. As we shall see soon, for a fixed choice of the normalization factor of the Haar measure μ for G , there is a natural choice for the normalization factor of the Haar measure $\hat{\mu}$ for \hat{G} .

Examples of locally compact abelian groups that play an important role in this dissertation are: the m -dimensional real vector space \mathbb{R}^m ; the m -fold torus $\mathbb{T}^m := \mathbb{R}^m / \mathbb{Z}^m$ and more generally \mathbb{R}^m / Λ for an arbitrary lattice Λ in \mathbb{R}^m ; and the finite ‘discretized torus’ group $\mathbb{D}^m := \frac{1}{q} \mathbb{Z}^m / \mathbb{Z}^m \subset \mathbb{T}^m$ (which is isomorphic to $\mathbb{Z}^m / q\mathbb{Z}^m$) for a positive integer q . Figure 2.1 below shows the corresponding dual groups as well as the respective (dual) Haar measures as used in Chapter 3 of this thesis.

In some cases it will be useful to identify the quotient groups $\mathbb{T}^m = \mathbb{R}^m / \mathbb{Z}^m$ and $\mathbb{D}^m = \frac{1}{q} \mathbb{Z}^m / \mathbb{Z}^m$ with the respective representing sets

$$\mathbb{T}_{\text{rep}}^m := [-\frac{1}{2}, \frac{1}{2})^m \subset \mathbb{R}^m \quad \text{and} \quad \mathbb{D}_{\text{rep}}^m := \frac{1}{q} \mathbb{Z}^m \cap \mathbb{T}_{\text{rep}}^m,$$

and similarly $\hat{\mathbb{D}}^m \simeq \mathbb{Z}^m / q\mathbb{Z}^m$ with

$$\hat{\mathbb{D}}_{\text{rep}}^m := [q]_c^m := \mathbb{Z}^m \cap [-\frac{q}{2}, \frac{q}{2})^m.$$

¹Discrete (and in particular, finite) groups have the discrete topology, implying that the continuity constraint for characters on these groups is void.

Group		Dual group	
G	μ	\hat{G}	$\hat{\mu}$
\mathbb{R}^m	λ	$\hat{\mathbb{R}}^m \simeq \mathbb{R}^m$	λ
$\mathbb{T}^m := \mathbb{R}^m / \mathbb{Z}^m$	λ	$\hat{\mathbb{T}}^m \simeq \mathbb{Z}^m$	$\#$
$\mathbb{D}^m := \frac{1}{q} \mathbb{Z}^m / \mathbb{Z}^m$	$\frac{1}{q^m} \#$	$\hat{\mathbb{D}}^m \simeq \mathbb{Z}^m / q\mathbb{Z}^m$	$\#$
\mathbb{R}^m / Λ	$\frac{1}{\det(\Lambda)} \lambda$	$\widehat{(\mathbb{R}^m / \Lambda)} \simeq \Lambda^*$	$\#$

Figure 2.1.: Some groups G and their respective dual groups \hat{G} , plus the considered (dual) Haar measures μ and $\hat{\mu}$. Here, λ denotes the Lebesgue measure and $\#$ the counting measure. Furthermore, Λ^* is the *dual lattice* of Λ , see Section 2.5.1.

It will be useful to understand that if $H \subset G$ is a closed subgroup then G/H and H have dual groups that satisfy the following natural isomorphisms.

$$\widehat{G/H} \simeq H^\perp := \{\chi \in \hat{G} \mid \chi(h) = 1 \text{ for all } h \in H\} \subset \hat{G} \quad \text{and} \quad \hat{H} \simeq \hat{G}/H^\perp.$$

As we shall see soon, for any choice of the Haar measure μ_H for H there is a natural choice for the Haar measure $\mu_{G/H}$ for G/H , and vice versa.

2.2.2. Norms and Fourier Transforms

Let G be as above with a fixed choice for the Haar measure μ . For any $p \in [1, \infty]$, $L_p(G)$ denotes the metric vector space of measurable functions $f : G \rightarrow \mathbb{C}$ with finite norm $\|f\|_p$ (modulo the functions with norm zero²), where

$$\|f\|_p^p := \int_{g \in G} |f(g)|^p d\mu \quad \text{for } p < \infty,$$

and

$$\|f\|_\infty := \text{ess sup}_{g \in G} |f(g)|,$$

²This in order to make $\|\cdot\|_p$ a metric: $\|f\|_p = 0$ implies $f = 0$ in that case.

the essential supremum of $|f|$. We write $\|f\|_{p,G}$ if we want to make G explicit. For any function $f \in L^1(G)$, the *Fourier transform* of f is the function

$$\mathcal{F}_G\{f\} : \hat{G} \rightarrow \mathbb{C}, \chi \mapsto \int_{g \in G} f(g) \bar{\chi}(g) d\mu,$$

also denoted by \hat{f} when G is clear from the context. The Fourier transform of $f \in L^1(G)$ is continuous, but not necessarily in $L^1(\hat{G})$.

For example, for the group $\mathbb{D}^m := \frac{1}{q} \mathbb{Z}^m / \mathbb{Z}^m$ with the Haar measure as fixed in Figure 2.1, the L_2 -norm and the Fourier transform are respectively given by

$$\|f\|_2^2 = \frac{1}{q^m} \sum_{x \in \mathbb{D}^m} |f(x)|^2 \quad \text{and} \quad \mathcal{F}\{f\}(y) = \frac{1}{q^m} \sum_{x \in \mathbb{D}^m} f(x) e^{-2\pi i \langle x, y \rangle}.$$

We note that we use a different convention on the scaling than what is common in the context of the quantum Fourier transform. Namely, in most literature (e.g., [NC11, §5.1]), the standard quantum Fourier transform uses a scaling of $q^{-m/2}$, for sake of preserving the L_2 -norm and symmetry; here, we use the scaling q^{-m} one way, and a unit scaling the other way.

Given the Haar measure μ for G , there exists a unique *dual* Haar measure $\hat{\mu}$ for \hat{G} with the property that, for any $f \in L^1(G)$, if $\hat{f} = \mathcal{F}_G\{f\} \in L^1(\hat{G})$, then $f = \mathcal{F}_G^{-1}\{\hat{f}\}$, where

$$\mathcal{F}_G^{-1}\{\hat{f}\} : G \rightarrow \mathbb{C}, g \mapsto \int_{\chi \in \hat{G}} \hat{f}(\chi) \chi(g) d\hat{\mu}$$

is the *inverse Fourier transform*. From now on it is always understood that the Haar measure of the dual group is chosen to be the dual of the Haar measure of the primal group. With this choice, we also have the following well known fact [DE16, Thm. 3.4.8].

Theorem 2.1 (Plancherel's Identity). *For all $f \in L^1(G) \cap L^2(G)$,*

$$\|f\|_{2,G} = \|\mathcal{F}_G\{f\}\|_{2,\hat{G}}.$$

Finally, we recall the *convolution theorem*, which states that $\widehat{fg} = \widehat{f} \star \widehat{g} = \int_{x \in G} \widehat{f}(x) \widehat{g}(\cdot - x) d\mu(x)$ for all functions $f, g \in L^1(G)$ that have Fourier transforms $\widehat{f}, \widehat{g} \in L^1(G)$. This extends to functions $f \in L^1(G/H)$ and $g \in L^1(G)$, with f understood as an H -periodic function on G . Tailored to $G = \mathbb{R}^m$ and $H = \Lambda$, where \mathbb{R}^m/Λ has dual group Λ^* , it then states that, for all $y \in \mathbb{R}^m$,

$$\begin{aligned} \mathcal{F}_{\mathbb{R}^m}\{fg\}(y) &= \mathcal{F}_{\mathbb{R}^m/\Lambda}\{f\} \star \mathcal{F}_{\mathbb{R}^m}\{g\}(y) \\ &= \sum_{\ell^* \in \Lambda^*} \mathcal{F}_{\mathbb{R}^m/\Lambda}\{f\}(\ell^*) \mathcal{F}_{\mathbb{R}^m}\{g\}(y - \ell^*). \end{aligned} \quad (2.4)$$

2.2.3. The Poisson Summation Formula

Poisson summation formula is well-known for the group $G = \mathbb{R}$, where it states that $\sum_{k \in \mathbb{Z}} \widehat{f}(k) = \sum_{x \in \mathbb{Z}} f(x)$. In the case $G = \mathbb{Z}/N\mathbb{Z}$, it reads

$$\sum_{i=0}^{N/s} \widehat{f}(is) = \sum_{j=1}^s f(j \frac{N}{s})$$

for any integer s that divides N . In order to formulate the Poisson summation formula for an arbitrary locally compact abelian group G , we need to introduce the notion of *restriction* and *periodization* of functions (see Figures 2.2 and 2.3).

Definition 2.2 (Restriction). *Let $H \subseteq G$ be a subset or a subgroup. For any continuous function $f : G \rightarrow \mathbb{C}$ we define $f|_H : H \rightarrow \mathbb{C}, h \mapsto f(h)$.*

Definition 2.3 (Periodization). *Let H be a closed subgroup of G with Haar measure μ_H . For any function $f \in L^1(G)$, we define*

$$f|^{G/H} : G/H \rightarrow \mathbb{C}, g + H \mapsto \int_{h \in H} f(g + h) d\mu_H.$$

A function on the real line \mathbb{R} ,



and its restriction to the integers \mathbb{Z}

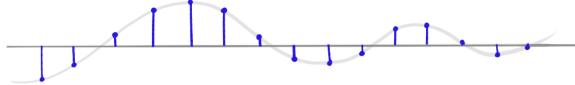


Figure 2.2.: A function on the real line and its restriction on the integers.

For any closed subgroup H of G with respective choices of Haar measures μ and μ_H , there exists a unique Haar measure $\mu_{G/H}$ for G/H such that the *quotient integral formula*

$$\int_{G/H} \left(\int_H f(g+h) d\mu_H(h) \right) d\mu_{G/H}(g+H) = \int_G f(g) d\mu(g)$$

holds for any continuous function $f : G \rightarrow \mathbb{C}$ with compact support (see [DE16, Sec. 1.5]).

With this choice of Haar measure for G/H , and with the dual measures for the respective dual groups, we are ready to state the general form of the Poisson summation formula (obtained from [DE16, Sec. 3.6], see also Figure 2.5).

Theorem 2.4 (Poisson Summation Formula). *For continuous $f \in L^1(G)$,*

$$\mathcal{F}_H\{f|_H\} = \mathcal{F}_G\{f\}|_{\widehat{H}} \quad \text{and} \quad \mathcal{F}_{G/H}\{f|^{G/H}\} = \mathcal{F}_G\{f\}|_{\widehat{G/H}}.$$

Applied to $G = \mathbb{R}^m$ and $H = \mathbb{Z}^m$, so that $G/H = \mathbb{R}^m/\mathbb{Z}^m = \mathbb{T}^m$ and $\widehat{G/H} \simeq \mathbb{Z}^m$; and applied to $G = \mathbb{T}^m$ and $H = \mathbb{D}^m$ below, we obtain the following.

Corollary 2.5. *For continuous $h \in L^1(\mathbb{R}^m)$, we have*

$$\mathcal{F}_{\mathbb{T}^m}\{h|_{\mathbb{T}^m}\} = \mathcal{F}_{\mathbb{R}^m}\{h\}|_{\mathbb{Z}^m}.$$

The real line \mathbb{R} ,

wrapped into
a helix,



made denser
and denser,

collapses into
the circle S^1 .

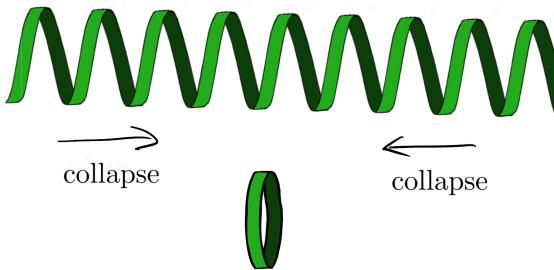


Figure 2.3.: The periodization of a function is a consequence of folding the space of its domain, i.e., taking the topological quotient space. In this example, the real line \mathbb{R} is folded into a circle.

Corollary 2.6. *For continuous $t \in L^1(\mathbb{T}^m)$, we have*

$$\mathcal{F}_{\mathbb{D}^m} \left\{ t|_{\mathbb{D}^m} \right\} = \mathcal{F}_{\mathbb{T}^m} \{ t \}|^{\hat{\mathbb{D}}^m}.$$

Remark 2.7. *The Poisson summation formula can be used to show that a ‘wide’ periodized Gaussian on the circle is close to a constant function, see Figure 2.7. The wider a Gaussian function, the narrower the Gaussian function of its Fourier transform is. Taking the restriction of such a ‘narrow’ Gaussian function to the integers \mathbb{Z} results in a spectrum heavily concentrated on zero, which corresponds to a constant function, as can be seen in the bottom example of Figure 2.7. Also note that for the ‘narrower’ Gaussian function on the circle, both the Gaussian on the circle as the restricted Fourier transform on \mathbb{Z} resemble much more a ‘real’ Gaussian function. In short, the narrower the Gaussian on the circle, the more Gaussian properties*

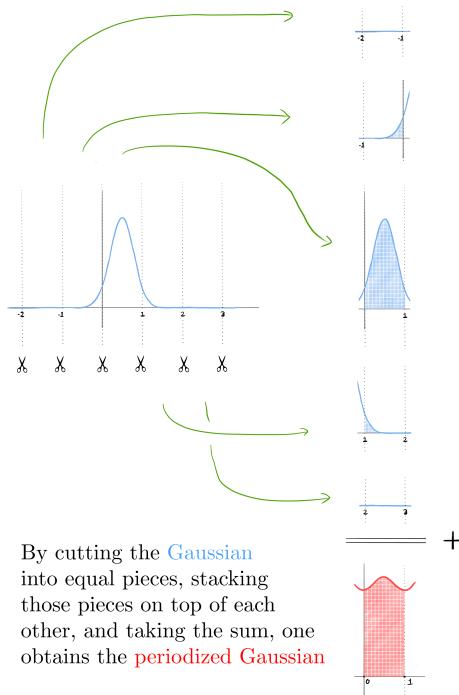


Figure 2.4.: An example of the periodization of a Gaussian on the real line, with respect to the subgroup $\mathbb{Z} \subseteq \mathbb{R}$. This leads to a *periodized Gaussian* on the circle $\mathbb{R}/\mathbb{Z} \simeq S^1$.

is has; the wider the Gaussian on the circle, the more ‘constant’ properties it has.

2.2.4. The Fourier Transform of Vector-valued Functions

The Fourier transform as discussed above generalizes to vector-valued functions $\mathbf{f} : G \rightarrow \mathbb{C}^N$ simply by applying \mathcal{F} to the N coordinate functions, resulting in a function $\mathcal{F}\{\mathbf{f}\} : \hat{G} \rightarrow \mathbb{C}^N$. By fixing an orthonormal basis, this extends to functions $\mathbf{f} : G \rightarrow \mathcal{H}$ for an arbitrary finite-dimensional complex Hilbert space, where, by linearity of the Fourier transform, $\mathcal{F}\{\mathbf{f}\} : \hat{G} \rightarrow \mathcal{H}$ is independent of the choice of the basis.

$$\begin{array}{ccccc}
 L^1(H) & \xleftarrow{|_H} & L^1(G) & \xrightarrow{|^{G/H}} & L^1(G/H) \\
 \mathcal{F}_H \downarrow & & \mathcal{F}_G \downarrow & & \mathcal{F}_{G/H} \downarrow \\
 L^1(\hat{G}/\widehat{G/H}) & \xleftarrow{|_{\hat{H}}} & L^1(\hat{G}) & \xrightarrow{|_{\widehat{G/H}}} & L^1(\widehat{G/H})
 \end{array}$$

Figure 2.5.: Informal illustration of Theorem 2.4 by means of a diagram that commutes whenever the maps are well defined.

The norm $\|\cdot\|_{2,G}$ on functions $G \rightarrow \mathbb{C}$ generalizes to vector-valued functions $\mathbf{f} : G \rightarrow \mathcal{H}$, as well, by defining $\|\mathbf{f}\|_{2,G}$ to be the norm of the scalar function $x \mapsto \|\mathbf{f}(x)\|_{\mathcal{H}} = \sqrt{\langle \mathbf{f}(x) | \mathbf{f}(x) \rangle}$. The vectorial Fourier transforms and norms are compatible with each other, in the sense that Plancherel's identity (see Theorem 2.1) still holds; that is,

$$\|\mathbf{f}\|_{2,G} = \|\mathcal{F}_G\{\mathbf{f}\}\|_{2,\hat{G}}. \quad (2.5)$$

Also the Poisson summation formula (see Theorem 2.4) is still valid, as well as the convolution theorem whenever one of the functions in the product is scalar:

$$\mathcal{F}_G\{\mathbf{f}g\} = \mathcal{F}_G\{\mathbf{f}\} \star \mathcal{F}_G\{g\}. \quad (2.6)$$

An important example is the case $\mathbf{f} : \mathbb{R}^m/\Lambda \rightarrow \mathcal{H}$. Spelling out the above, we get

$$\begin{aligned}
 \mathcal{F}_{\mathbb{R}^m/\Lambda}\{\mathbf{f}\} : \Lambda^* &\rightarrow \mathcal{H}, \\
 \ell^* &\mapsto |c_{\ell^*}\rangle := \frac{1}{\det \Lambda} \int_{x \in \mathbb{R}^m/\Lambda} |\mathbf{f}(x)\rangle e^{-2\pi i \langle x, \ell^* \rangle} dx,
 \end{aligned} \quad (2.7)$$

where the vectors $|c_{\ell^*}\rangle$ are also referred to as the *(vectorial) Fourier coefficients* of \mathbf{f} . The Parseval-Plancherel identity [DE16, Thm. 3.4.8] then becomes

$$\sum_{\ell^* \in \Lambda^*} \langle c_{\ell^*} | c_{\ell^*} \rangle = \|\mathbf{f}\|_{\mathbb{R}^m/\Lambda}^2 := \frac{1}{\det \Lambda} \int_{x \in \mathbb{R}^m/\Lambda} \langle \mathbf{f}(x) | \mathbf{f}(x) \rangle dx. \quad (2.8)$$

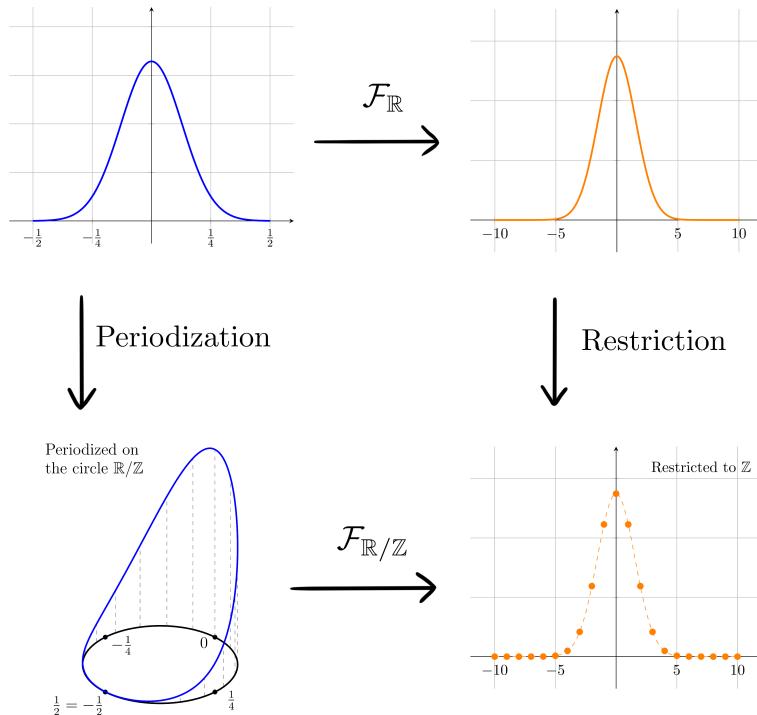


Figure 2.6.: A graphical depiction of the Poisson summation formula as described in Theorem 2.4, applied to a Gaussian function. First periodizing a function and then applying the Fourier transform gives the same result as first applying the Fourier transform and then restricting the function. As a result, the Fourier transform of a *periodized Gaussian* is a *discrete Gaussian*.

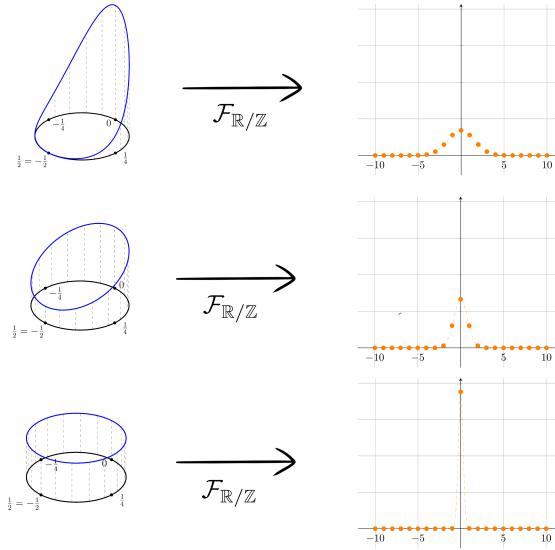


Figure 2.7.: The narrower the Gaussian on the circle, the more it looks like a Gaussian; the wider the Gaussian on the circle, the closer to constant it is.

The convolution theorem, as in Equation (2.6) and Equation (2.4), in this case, becomes,

$$\begin{aligned} \mathcal{F}_{\mathbb{R}^m} \{ \mathbf{f} g \} &= \mathcal{F}_{\mathbb{R}^m / \Lambda} \{ \mathbf{f} \} \star \mathcal{F}_{\mathbb{R}^m} \{ g \} \\ &= \sum_{\ell^* \in \Lambda^*} \mathcal{F}_{\mathbb{R}^m / \Lambda} \{ \mathbf{f} \} \cdot \mathcal{F}_{\mathbb{R}^m} \{ g \}(\cdot - \ell^*). \end{aligned} \quad (2.9)$$

2.2.5. Trigonometric Approximation

As another application of the Poisson summation formula, we derive a relation between the Lipschitz constant of a function on $\mathbb{T}^m = \mathbb{R}^m / \mathbb{Z}^m$ and the ‘error of discretization’ in the Fourier transform when restricting the function to \mathbb{D}^m .

Theorem 2.8. *For any Lipschitz function $\mathbf{h} : \mathbb{T}^m \rightarrow \mathcal{H}$ (where \mathcal{H} is a Hilbert space) with Lipschitz constant $\text{Lip}(\mathbf{h})$, and any subset $C \subseteq \mathbb{D}^m$, we*

have

$$\left| \|1_C \cdot \mathcal{F}_{\mathbb{D}^m} \{\mathbf{h}\}\|_{\hat{\mathbb{D}}^m} - \|1_C \cdot \mathcal{F}_{\mathbb{T}^m} \{\mathbf{h}\}\|_{\mathbb{Z}^m} \right| \leq \frac{4\pi\sqrt{m} \cdot \text{Lip}(\mathbf{h})}{q}.$$

Here and below, we slightly abuse notation and use 1_C as indicator function acting on $\hat{\mathbb{D}}^m$ and on \mathbb{Z}^m , justified by identifying $\hat{\mathbb{D}}^m$ with $\hat{\mathbb{D}}_{\text{rep}}^m = [q]_c^m \subset \mathbb{Z}^m$. Also, we write $\mathcal{F}_{\mathbb{D}^m} \{\mathbf{h}\}$ instead of $\mathcal{F}_{\mathbb{D}^m} \{\mathbf{h}|_{\mathbb{D}^m}\}$, taking it as understood that \mathbf{h} is restricted to \mathbb{D}^m when applying $\mathcal{F}_{\mathbb{D}^m}$.

Proof. Using a result of Yudin ([Yud76, Example I after Thm. 2], see also³ Appendix A.4), there exists a trigonometric approximation \mathbf{t} of \mathbf{h} , i.e. a function $\mathbf{t} : \mathbb{T}^m \rightarrow \mathbb{C}$ with $\hat{\mathbf{t}}(x) := \mathcal{F}_{\mathbb{T}^m} \{\mathbf{t}\}(x) = 0$ for all $x \in \mathbb{Z}^m \setminus [q]_c^m$ so that $\|\mathbf{h} - \mathbf{t}\|_\infty \leq \pi\sqrt{m} \cdot \text{Lip}(\mathbf{h})/q$. Recalling that $\hat{\mathbb{D}}^m \simeq \mathbb{Z}^m/q\mathbb{Z}^m$, the fact that $\hat{\mathbf{t}} : \mathbb{Z}^m \rightarrow \mathbb{C}$ vanishes outside of $[q]_c^m$ implies for all $x \in [q]_c^m$ that

$$\hat{\mathbf{t}}(x) = \sum_{d \in q\mathbb{Z}^m} \hat{\mathbf{t}}(x + d) = \hat{\mathbf{t}}|_{\hat{\mathbb{D}}^m}(x) = \mathcal{F}_{\mathbb{D}^m} \{\mathbf{t}\}(x),$$

where the last equality holds by Corollary 2.6 (and our convention of omitting the restriction to \mathbb{D}^m). In particular, we have $\|1_C \cdot \mathcal{F}_{\mathbb{D}^m} \{\mathbf{t}\}\|_{2, \hat{\mathbb{D}}^m} = \|1_C \cdot \mathcal{F}_{\mathbb{T}^m} \{\mathbf{t}\}\|_{2, \mathbb{Z}^m}$. Therefore, by the (reverse) triangle inequality and the linearity of the Fourier transform, one obtains

$$\begin{aligned} & \left| \|1_C \cdot \mathcal{F}_{\mathbb{D}^m} \{\mathbf{h}\}\|_{\hat{\mathbb{D}}^m} - \|1_C \cdot \mathcal{F}_{\mathbb{T}^m} \{\mathbf{h}\}\|_{\mathbb{Z}^m} \right| \\ & \leq \|1_C \cdot \mathcal{F}_{\mathbb{D}^m} \{\mathbf{h} - \mathbf{t}\}\|_{\hat{\mathbb{D}}^m} + \|1_C \cdot \mathcal{F}_{\mathbb{T}^m} \{\mathbf{h} - \mathbf{t}\}\|_{\mathbb{Z}^m}. \end{aligned}$$

We now observe that

$$\begin{aligned} \|1_C \cdot \mathcal{F}_G \{\mathbf{h} - \mathbf{t}\}\|_{2, \hat{G}} & \leq \|\mathcal{F}_G \{\mathbf{h} - \mathbf{t}\}\|_{2, \hat{G}} = \|\mathbf{h} - \mathbf{t}\|_{2, G} \leq \sqrt{\mu(G)} \|\mathbf{h} - \mathbf{t}\|_\infty \\ & \leq \sqrt{\mu(G)} \cdot \pi\sqrt{m} \cdot \text{Lip}(\mathbf{h})/q, \end{aligned}$$

where $\mu(G) = \int_G d\mu$ denotes the total measure of G . We conclude by noting that $\mu(G) = 1$ for both groups at hand $G = \mathbb{D}^m$ and $G = \mathbb{T}^m$. \square

³In Appendix A.4, we provide a slight generalization of Yudin's paper [Yud76] to functions with vectorial output. In principle the bound of Theorem 2.8 can also be derived without this generalization, but at the cost of an undesirable extra factor $\dim \mathcal{H} = 2^n$.

2.3. Number Theory

2.3.1. Algebraic Number Theory

In this thesis is assumed that the reader is somewhat familiar with the main concepts of algebraic number theory. In this section, we very briefly introduce definitions and notions required for this thesis. For a more elaborate explanation, I would suggest Neukirch's textbook [NS13].

Throughout this thesis, we use a fixed number field K of degree $n \geq 3$ over \mathbb{Q} , having ring of integers \mathcal{O}_K , discriminant Δ_K , regulator R_K , class number h_K and group of roots of unity μ_K . Elements of the number field K are generally denoted by lowercase Greek letters, α, β, γ , etc. Minkowski's theorem [Min67, p. 261–264] states⁴ that $\log |\Delta_K| \geq \log(2) \cdot n$. The number field K has n field embeddings into \mathbb{C} , which are divided in $n_{\mathbb{R}}$ real embeddings and $n_{\mathbb{C}}$ conjugate pairs of complex embeddings, i.e., $n = n_{\mathbb{R}} + 2n_{\mathbb{C}}$. These embeddings combined yield the so-called Minkowski embedding $K \rightarrow K_{\mathbb{R}} \subseteq \bigoplus_{\sigma: K \hookrightarrow \mathbb{C}} \mathbb{C}$, $\alpha \mapsto (\sigma(\alpha))_{\sigma}$, where

$$K_{\mathbb{R}} = \left\{ x \in \bigoplus_{\sigma: K \hookrightarrow \mathbb{C}} \mathbb{C} \mid x_{\bar{\sigma}} = \overline{x_{\sigma}} \right\}.$$

Here, $\bar{\sigma}$ equals the conjugate embedding of σ whenever σ is a complex embedding and it is just σ itself whenever it is a real embedding. Note that we index the components of the vectors in $K_{\mathbb{R}}$ by the embeddings of K . Embeddings up to conjugation are called infinite places, denoted by ν . With any embedding σ we denote by ν_{σ} the associated place; and for any place ν we choose a fixed embedding σ_{ν} . There are also *finite* places ν , which are in one-to-one correspondence with the prime ideals of \mathcal{O}_K . For finite places $\nu \nmid \infty$ we denote by $\mathfrak{p}_{\nu} \in \mathcal{I}_K$ their associated prime ideal, for infinite places $\nu \mid \infty$ we denote by σ_{ν} their (chosen) associated embedding.

Composing the Minkowski embedding by the component-wise logarithm of

⁴By Minkowski's theorem, we have $|\Delta_K|^{1/n} \geq \pi/4 \cdot \frac{n^2}{(n!)^{2/n}} \geq 2$ for $n \geq 3$.

the entries' absolute values yields the logarithmic map, denoted by Log .

$$\text{Log} : K^* \rightarrow \text{Log } K_{\mathbb{R}} \subseteq \bigoplus_{\sigma: K \hookrightarrow \mathbb{C}} \mathbb{R}, \alpha \mapsto (\log |\sigma(\alpha)|)_{\sigma}.$$

The multiplicative group of integral units \mathcal{O}_K^{\times} under the logarithmic map forms a lattice, namely the lattice $\Lambda_K = \text{Log}(\mathcal{O}_K^{\times}) \subseteq \text{Log } K_{\mathbb{R}}$ (see Section 2.5.1 for the preliminaries on lattices). This so-called logarithmic unit lattice has rank $r = n_{\mathbb{R}} + n_{\mathbb{C}} - 1$, is orthogonal to the all-one vector $(1)_{\sigma}$, and has covolume $\text{Vol}(\Lambda_K) = \sqrt{n} \cdot 2^{-n_{\mathbb{C}}/2} \cdot R_K$, where the $2^{-n_{\mathbb{C}}/2}$ factor is due to the specific embedding we use (see Lemma A.2). We denote by $H = \text{Span}(\Lambda_K)$ the hyperplane of dimension r , which can also be defined as the subspace $\text{Log}(K_{\mathbb{R}}^0)$ of $\text{Log } K_{\mathbb{R}}$, where

$$K_{\mathbb{R}}^0 = \{x \in K_{\mathbb{R}} \mid \prod_{\sigma: K \hookrightarrow \mathbb{C}} |x_{\sigma}| = 1\}.$$

In other words, $H = \text{Log } K_{\mathbb{R}}^0$ is the subspace of $\text{Log } K_{\mathbb{R}}$ orthogonal to the all-one vector $(1)_{\sigma}$. We denote by $T = H/\Lambda_K$ the hypertorus defined by the logarithmic unit lattice Λ_K . Note that $K_{\mathbb{R}} \simeq \prod_{\nu} K_{\nu}$, where ν ranges over all infinite places of K , and $K_{\nu} = \mathbb{C}$ or \mathbb{R} depending on whether ν is complex or real respectively. In some cases it is more convenient to use this particular viewpoint of $K_{\mathbb{R}}$. Note that $K_{\mathbb{R}}^0$ can then be identified with

$$K_{\mathbb{R}}^0 = \{x \in \prod_{\nu \mid \infty} K_{\nu} \mid \prod_{\nu \mid \infty} |x_{\nu}|_{\mathbb{C}}^{[K_{\nu} : \mathbb{R}]} = 1\}. \quad (2.10)$$

Note that we take the usual complex absolute value here, which is raised to the power two whenever $K_{\nu} = \mathbb{C}$ and to the power one otherwise.

Fractional ideals of the number field K are denoted by $\mathfrak{a}, \mathfrak{b}, \dots$, but the symbols $\mathfrak{p}, \mathfrak{q}$ are generally reserved for integral prime ideals of \mathcal{O}_K . Also, the symbol \mathfrak{m} is reserved for the modulus ideal $\mathfrak{m} \subseteq \mathcal{O}_K$, a notion from class field theory. One can think of the primes dividing \mathfrak{m} as the primes 'to avoid'. For $\mathfrak{a} \in \mathcal{I}_K$, we denote $\text{ord}_{\mathfrak{p}}(\mathfrak{a}) = \max\{k \mid \mathfrak{p}^k \text{ divides } \mathfrak{a}\}$ for the \mathfrak{p} -valuation of the ideal \mathfrak{a} ; this can be generalized for elements $\alpha \in K^*$ by considering the principal ideal generated by that element. The group of fractional ideals of K is denoted by \mathcal{I}_K ; the group of fractional ideals coprime with \mathfrak{m} is

denoted by $\mathcal{I}_K^{\mathfrak{m}}$. Principal ideals with generator $\alpha \in K^*$ are usually denoted by (α) . We denote by $K^{\mathfrak{m},1} = \langle \alpha \in \mathcal{O}_K \mid \alpha \equiv 1 \pmod{\mathfrak{m}} \rangle$ the *ray* modulo \mathfrak{m} , i.e., the multiplicative subgroup of K^* generated by elements in \mathcal{O}_K that are one modulo \mathfrak{m} . In many texts the modulus can also include infinite primes (i.e., embeddings into \mathbb{C}); not in this thesis.

For any integral ideal \mathfrak{a} , we define the norm $\mathcal{N}(\mathfrak{a})$ of \mathfrak{a} to be the number $|\mathcal{O}_K/\mathfrak{a}|$; this norm then generalizes to fractional ideals and elements as well. The class group of \mathcal{O}_K , denoted by Cl_K , is the quotient of the group \mathcal{I}_K by the subgroup of principal ideals $\text{Princ}_K := \{(\alpha) \in \mathcal{I}_K \mid \alpha \in K\}$. For any fractional ideal \mathfrak{a} , we denote the ideal class of \mathfrak{a} in Cl_K by $[\mathfrak{a}]$.

In some parts of this thesis we need the notion of the *idèle group* \mathcal{J}_K , which is a topological group defined by the restricted topological product of the completions of the number field K over all places $\prod_{\nu} K_{\nu}^*$ where the restriction is with respect to the unit groups $\mathcal{O}_{\nu}^{\times} \subseteq K_{\nu}^*$. For a modulus \mathfrak{m} , the idèle group modulo \mathfrak{m} , $\mathcal{J}_{K^{\mathfrak{m}}}$, is defined similarly, by just leaving out the completions whose place are associated with a prime dividing \mathfrak{m} . For any modulus \mathfrak{m} , the ray $K^{\mathfrak{m},1}$ embeds diagonally into $\mathcal{J}_{K^{\mathfrak{m}}}$, by $\alpha \mapsto (\alpha_{\nu})_{\nu} \in \mathcal{J}_{K^{\mathfrak{m}}}$. Each component of this diagonal map is just the embedding of the completion $K \rightarrow K_{\nu}$. The quotient of the idèle group (modulo \mathfrak{m}) and the ray is called the *idèle class group* \mathcal{C}_K , which can be shown to be the same for any modulus \mathfrak{m} (see [Lan12, Ch. VII, §4]).

In this thesis, extra attention is paid to the cyclotomic number fields $K = \mathbb{Q}(\zeta_m)$, for which one can sometimes phrase sharper results due to the fact thats these fields have more structure. The result in Chapter 5 tailored to cyclotomic fields relies on the size of the class group $h_K^+ = |\text{Cl}_{K^+}|$ of the maximum real subfield $K^+ = \mathbb{Q}(\zeta_m + \bar{\zeta}_m)$ of K , which is conjectured to be rather small [Mil15; BPR04]. In Chapter 5, we make the mild assumption that $h_K^+ \leq (\log n)^{O(n)}$, where $n = [K : \mathbb{Q}] = \phi(m)$.

An important identity that will play a large role throughout this thesis is the *class number formula*, which relates multiple number-theoretic quantities with the residue at $s = 1$ of the Dedekind zeta function $\zeta_K(s) = \sum_{\mathfrak{a} \subseteq \mathcal{O}_K} \frac{1}{\mathcal{N}(\mathfrak{a})^s}$.

$$\rho_K = \lim_{s \rightarrow 1} (s-1) \zeta_K(s) \frac{2^{n_{\mathbb{R}}} \cdot (2\pi)^{n_{\mathbb{C}}} \cdot R_K \cdot h_K}{|\mu_K| \cdot \sqrt{|\Delta_K|}}. \quad (2.11)$$

2.3.2. The Extended Riemann Hypothesis

Almost all results in this paper rely heavily on the *Extended Riemann Hypothesis* (in the subsequent part of this paper abbreviated by ERH), which refers to the Riemann Hypothesis extended to Hecke L -functions (see [IKS04, §5.7]). All statements that mention (ERH), such as Theorem 4.3, assume the Extended Riemann Hypothesis.

Definition 2.9 (Hecke L-function). *Let K be a number field and let $\chi : \mathcal{J}_{K^{\mathfrak{m}}}/K^{\mathfrak{m},1} \rightarrow S^1$ be a Hecke character on the idèle class group $\mathcal{C}_K = \mathcal{J}_{K^{\mathfrak{m}}}/K^{\mathfrak{m},1}$ of K (see [NS13, Ch. VI and Ch. VII, §6] and Section 4.3.4) defined modulo its conductor \mathfrak{m} . Then we define*

$$L(\chi, s) = \sum_{\substack{\mathfrak{a} \subseteq \mathcal{O}_K \\ \mathfrak{a} + \mathfrak{m} = \mathcal{O}_K}} \frac{\chi(\mathfrak{a})}{\mathcal{N}(\mathfrak{a})^s}$$

to be the associated Hecke L -function, where the sum ranges over all integral ideals of the maximal order \mathcal{O}_K of K , coprime with the modulus \mathfrak{m} (see, for example [Neu85, Ch. V, Def. 3.1]).

Definition 2.10 (Extended Riemann Hypothesis). *For all number fields K and all Hecke characters χ , all zeroes of the Hecke L -functions that are in the critical strip $0 < \operatorname{Re}(s) < 1$, satisfy $\operatorname{Re}(s) = 1/2$. I.e., for all number fields K , Hecke characters χ and all complex numbers $s \in \mathbb{C}$,*

$$[L(\chi, s) = 0 \text{ and } \operatorname{Re}(s) \in (0, 1)] \implies \operatorname{Re}(s) = 1/2.$$

Remark 2.11. *Most of the results in this thesis are phrased in terms of a fixed number field K . In such a case it is of course not needed to assume the Extended Riemann Hypothesis for all number fields; it suffices*

to assume the Extended Riemann Hypothesis for Hecke L -functions arising from Hecke-characters for the fixed number field K .

So, if a theorem in this thesis regards only a single number field K , and it assumes the Extended Riemann Hypothesis, one may weaken this hypothesis to the Extended Riemann Hypothesis ‘tailored to K ’.

2.3.3. Prime Densities

In multiple parts of this paper, we need an estimate on the number of prime ideals with bounded norm. This is achieved in the following theorem, obtained from Bach and Shallit’s book [BS96, Thm. 8.7.4].

Theorem 2.12 (ERH). *Let $\pi_K(x)$ be the number of prime ideals of K of norm $\leq x$. Then, assuming the Extended Riemann Hypothesis, there exists an absolute constant C (i.e., independent of K and x) such that, for all $x \geq 2$,*

$$|\pi_K(x) - \text{li}(x)| \leq C \cdot \sqrt{x} (n \log x + \log |\Delta_K|),$$

where $\text{li}(x) = \int_2^x \frac{dt}{\ln t} \sim \frac{x}{\ln x}$.

In certain cases, we prefer a more explicit variant of this theorem that is due to Grenié and Molteni [GM15, Cor. 1.4].

Lemma 2.13 (ERH). *Let $\mathfrak{m} \subseteq \mathcal{O}_K$ be an ideal modulus and denote*

$$\pi_K^{\mathfrak{m}}(x) = |\{\mathfrak{p} \in \mathcal{I}_K^{\mathfrak{m}} \mid \mathfrak{p} \text{ prime and } \mathcal{N}(\mathfrak{p}) \leq x\}|$$

for the number of prime ideals not dividing \mathfrak{m} and having norm bounded by $x \in \mathbb{R}$. Let $\omega(\mathfrak{m})$ denote the number of different prime ideal divisors of \mathfrak{m} .

Then, for all $x \geq \max((12 \log |\Delta_K| + 8n + 28)^4, 3 \cdot 10^{11}, 16 \cdot \omega(\mathfrak{m})^2)$, we have

$$\pi_K^{\mathfrak{m}}(x) \geq \frac{x}{4 \ln x}.$$

Proof. Denote $\pi_K(x) = |\{\mathfrak{p} \in \mathcal{I}_K \mid \mathfrak{p} \text{ prime and } \mathcal{N}(\mathfrak{p}) \leq x\}|$, i.e., whenever $\mathfrak{m} = \mathcal{O}_K$. We will prove the statement for this specific case first. By simplifying an explicit result of Grenié and Molteni [GM15, Cor. 1.4], we obtain, under the Extended Riemann Hypothesis⁵,

$$\left| \pi_K(x) - \pi_K(3) - \int_3^x \frac{du}{\log u} \right| \leq \sqrt{x}[6 \log |\Delta_K| + 4n \log x + 14].$$

Therefore, we have

$$\begin{aligned} \pi_K(x) &\geq \int_3^x \frac{du}{\log u} - \sqrt{x}[6 \log |\Delta_K| + 4n \log x + 14] \\ &\geq \frac{x}{\ln x} - \sqrt{x} \ln(x)[6 \log |\Delta_K| + 4n + 14] \\ &= \frac{x}{\ln x} \left(1 - \frac{\ln(x)^2(6 \log |\Delta_K| + 4n + 14)}{\sqrt{x}} \right) \geq \frac{x}{2 \ln x} \end{aligned}$$

where the first inequality follows from omitting $\pi_K(3)$ and the second inequality from $\int_3^x \frac{du}{\ln u} \geq \frac{x}{\ln x}$ and from the assumption that $x > 2^4 \cdot (6 \log |\Delta_K| + 4n + 14)^4$ and $x > 3 \cdot 10^{11}$. Note that with such x , we have $\ln(x)^2/\sqrt{x} < x^{-1/4}$, so that $\frac{\ln(x)^2(6 \log |\Delta_K| + 4n + 14)}{\sqrt{x}} < 1/2$.

For the general case of $\mathfrak{m} \neq \mathcal{O}_K$, we need to avoid \mathfrak{m} ; so writing $\omega(\mathfrak{m})$ for the number of different prime ideals dividing \mathfrak{m} , we obtain

$$\pi_K^{\mathfrak{m}}(x) \geq \pi_K(x) - \omega(\mathfrak{m}) \geq \frac{x}{2 \ln x} \left(1 - \frac{2 \cdot \omega(\mathfrak{m}) \cdot \ln x}{x} \right) \geq \frac{x}{4 \ln x}.$$

Where the last inequality can be deduced as follows. Since $x > 3 \cdot 10^{11}$, surely $\frac{\ln x}{x} \leq x^{-1/2} \leq (4 \cdot \omega(\mathfrak{m}))^{-1}$ and therefore $\frac{2 \cdot \omega(\mathfrak{m}) \cdot \ln x}{x} \leq 1/2$. This proves the claim. \square

Lemma 2.14 (Sampling of prime ideals, ERH). *Let a basis of \mathcal{O}_K be known and let $\mathcal{P} = \{\mathfrak{p} \text{ prime ideal of } K \mid \mathcal{N}(\mathfrak{p}) \leq B\}$ be the set of prime ideals of norm bounded by $B \geq \max((12 \log |\Delta_K| + 8n + 28)^4, 3 \cdot 10^{11})$. Then one can sample uniformly from \mathcal{P} in expected time $O(n^3 \log^2 B)$.*

⁵In the paper of Grenié and Molteni [GM15, Cor. 1.4], only the Dedekind zeta function $\zeta_K(s) = \sum_{\mathfrak{a}} \mathcal{N}(\mathfrak{a})^{-s}$ needs to satisfy the condition that all of its non-trivial zeroes lie at the vertical line $\Re(s) = 1/2$.

Proof. The sampling algorithm can be described as follows. Sample an integer uniformly in $[0, B]$ and check if it is a prime. If it is, factor the obtained prime p in \mathcal{O}_K and list the different prime ideal factors $\{\mathfrak{p}_1, \dots, \mathfrak{p}_k\}$ that have norm bounded by B . Choose one \mathfrak{p}_i uniformly at random in $\{\mathfrak{p}_1, \dots, \mathfrak{p}_k\}$ and output it with probability k/n . Otherwise, output ‘failure’.

Let $\mathfrak{q} \in \mathcal{P}$ be arbitrary, and let $\mathcal{N}(\mathfrak{q}) = q^j$ with q prime. Then, the probability of sampling \mathfrak{q} equals $\frac{1}{nB}$, namely $\frac{1}{n}$ times the probability of sampling q . Therefore, the probability of sampling successfully (i.e., no failure) equals $\frac{|\mathcal{P}|}{nB} \geq \frac{1}{2n \log B}$, since $|\mathcal{P}| \geq \frac{B}{2 \log B}$, by Lemma 2.13.

The most costly part of the algorithm is the factorization of a rational prime $p \leq B$ into prime ideals of \mathcal{O}_K . This can be performed using the Kummer-Dedekind algorithm, which essentially amounts to factoring a degree n polynomial modulo p . Using Shoup’s algorithm [Sho95] (which has complexity $O(n^2 + n \log p)$ [GP01, §4.1]) yields the complexity claim. \square

2.4. Arakelov Theory

2.4.1. The Arakelov Ray Divisor Group

The *Arakelov ray divisor group* with respect to a modulus $\mathfrak{m} \subseteq \mathcal{O}_K$ is the group

$$\text{Div}_{K^{\mathfrak{m}}} = \bigoplus_{\mathfrak{p} \nmid \mathfrak{m}} \mathbb{Z} \times \bigoplus_{\nu} \mathbb{R}$$

where \mathfrak{p} ranges over the set of all prime ideals of \mathcal{O}_K that do not divide the modulus \mathfrak{m} , and ν over the set of infinite primes (embeddings into the complex numbers up to possible conjugation). For readers that are not yet familiar with Arakelov ray divisor groups it might be insightful to first consider the ordinary Arakelov divisor group, which is obtained by putting $\mathfrak{m} = \mathcal{O}_K$.

We write an arbitrary element in $\text{Div}_{K^{\mathfrak{m}}}$ as

$$\mathbf{a} = \sum_{\mathfrak{p} \nmid \mathfrak{m}} n_{\mathfrak{p}} \cdot (\mathfrak{p}) + \sum_{\nu} x_{\nu} \cdot (\nu),$$

with only finitely many non-zero $n_{\mathfrak{p}}$. We will consistently use the symbols $\mathbf{a}, \mathbf{b}, \mathbf{e}, \dots$ for Arakelov ray divisors. Denoting $\text{ord}_{\mathfrak{p}}$ for the valuation at the prime \mathfrak{p} , there is a canonical homomorphism

$$\langle \cdot \rangle : K^{\mathfrak{m},1} \rightarrow \text{Div}_{K^{\mathfrak{m}}}, \quad \alpha \longmapsto \sum_{\mathfrak{p} \nmid \mathfrak{m}} \text{ord}_{\mathfrak{p}}(\alpha) \langle \mathfrak{p} \rangle - \sum_{\nu} \log |\sigma_{\nu}(\alpha)| \cdot \langle \nu \rangle.$$

The divisors of the form $\langle \alpha \rangle$ for $\alpha \in K^{\mathfrak{m},1}$ are called *principal ray divisors*. Here, $K^{\mathfrak{m},1} = \langle \alpha \in \mathcal{O}_K \mid \alpha \equiv 1 \pmod{\mathfrak{m}} \rangle$ is the multiplicative subgroup of K^* generated by elements equivalent to one modulo \mathfrak{m} . We will also make use of the notation $K^{\mathfrak{m}} = \langle \alpha \in \mathcal{O}_K \mid \alpha \pmod{\mathfrak{m}} \in (\mathcal{O}_K/\mathfrak{m})^* \rangle$, the multiplicative subgroup of K^* generated by elements coprime to \mathfrak{m} . Note that $K^{\mathfrak{m},1} \subseteq K^{\mathfrak{m}}$.

Just as the ideal ray class group is the group of ideals coprime with \mathfrak{m} quotiented by the ‘ray’ $K^{\mathfrak{m},1}$, the *Picard ray group* is the group of Arakelov ray divisors quotiented by the group of principal ray Arakelov divisors. In other words, the Picard ray group $\text{Pic}_{K^{\mathfrak{m}}}$ is defined by the following exact sequence, where $\mu_{K^{\mathfrak{m},1}} = \mu_K \cap K^{\mathfrak{m},1}$, the roots of unity in the ray.

$$0 \rightarrow \mu_{K^{\mathfrak{m},1}} \rightarrow K^{\mathfrak{m},1} \xrightarrow{\langle \cdot \rangle} \text{Div}_{K^{\mathfrak{m}}} \rightarrow \text{Pic}_{K^{\mathfrak{m}}} \rightarrow 0.$$

For any Arakelov ray divisor $\mathbf{a} = \sum_{\mathfrak{p} \nmid \mathfrak{m}} n_{\mathfrak{p}} \cdot \langle \mathfrak{p} \rangle + \sum_{\nu} x_{\nu} \cdot \langle \nu \rangle$, we denote its class in the Picard ray group $\text{Pic}_{K^{\mathfrak{m}}}$ by $[\mathbf{a}]$; in the same fashion that $[\mathfrak{a}]$ denotes the ideal class of the ideal \mathfrak{a} .

2.4.2. The Arakelov Ray Class Group

Despite the Arakelov ray divisor group and Picard ray group being interesting groups, it is for our purposes more useful to consider the *degree-zero* subgroups of these groups. The degree map is defined as follows:

$$\begin{aligned} \deg : \text{Div}_{K^{\mathfrak{m}}} &\rightarrow \mathbb{R}, \\ \sum_{\mathfrak{p} \nmid \mathfrak{m}} n_{\mathfrak{p}} \cdot \langle \mathfrak{p} \rangle + \sum_{\nu} x_{\nu} \cdot \langle \nu \rangle &\mapsto \sum_{\mathfrak{p} \nmid \mathfrak{m}} n_{\mathfrak{p}} \cdot \log(\mathcal{N}(\mathfrak{p})) + \sum_{\nu \text{ real}} x_{\nu} + \sum_{\nu \text{ complex}} 2 \cdot x_{\nu}. \end{aligned} \tag{2.12}$$

The degree map sends principal ray divisors (α) for $\alpha \in K^{\mathfrak{m},1}$ to zero; therefore, the degree map is properly defined on $\text{Pic}_{K^{\mathfrak{m}}}$, as well. We subsequently define the *degree-zero Arakelov ray divisor group* $\text{Div}_{K^{\mathfrak{m}}}^0 = \{\mathbf{a} \in \text{Div}_{K^{\mathfrak{m}}} \mid \deg(\mathbf{a}) = 0\}$ and the *Arakelov ray class group* $\text{Pic}_{K^{\mathfrak{m}}}^0 = \{[\mathbf{a}] \in \text{Pic}_{K^{\mathfrak{m}}} \mid \deg([\mathbf{a}]) = 0\}$. In other words, the group consisting of the degree zero Picard ray classes is called the Arakelov ray class group.

Any Arakelov ray divisor $\mathbf{a} \in \text{Div}_{K^{\mathfrak{m}}}^0$ can be decomposed in a finite and an infinite part, $\mathbf{a} = \mathbf{a}_f + \mathbf{a}_{\infty}$.

$$\mathbf{a} = \underbrace{\sum_{\mathfrak{p} \nmid \mathfrak{m}} n_{\mathfrak{p}} \cdot (\mathfrak{p})}_{\mathbf{a}_f} + \underbrace{\sum_{\nu} x_{\nu} \cdot (\nu)}_{\mathbf{a}_{\infty}} \quad (2.13)$$

The finite part \mathbf{a}_f , that consists of a formal integer sum of prime ideals, can be uniquely associated with an ideal in $\mathcal{I}_K^{\mathfrak{m}}$, i.e., we have

$$\text{Exp}(\cdot_f) : \text{Div}_{K^{\mathfrak{m}}}^0 \rightarrow \mathcal{I}_K^{\mathfrak{m}}, \quad \mathbf{a} \mapsto \text{Exp}(\mathbf{a}_f) = \prod_{\mathfrak{p} \nmid \mathfrak{m}} \mathfrak{p}^{n_{\mathfrak{p}}},$$

where we use the exponential function Exp to denote the map sending $\sum_{\mathfrak{p} \nmid \mathfrak{m}} n_{\mathfrak{p}} (\mathfrak{p})$ to $\prod_{\mathfrak{p} \nmid \mathfrak{m}} \mathfrak{p}^{n_{\mathfrak{p}}}$. This map $\text{Exp}(\cdot_f) : \text{Div}_{K^{\mathfrak{m}}}^0 \rightarrow \mathcal{I}_K^{\mathfrak{m}}$ has the hyperplane H as kernel via the inclusion $H \hookrightarrow \text{Div}_{K^{\mathfrak{m}}}^0$ and admits a section $d^0 : \mathcal{I}_K^{\mathfrak{m}} \rightarrow \text{Div}_{K^{\mathfrak{m}}}^0$, defined by the following rule.

$$d^0 : \mathcal{I}_K^{\mathfrak{m}} \rightarrow \text{Div}_{K^{\mathfrak{m}}}^0, \quad \mathbf{a} \mapsto \sum_{\mathfrak{p} \nmid \mathfrak{m}} \text{ord}_{\mathfrak{p}}(\mathbf{a}) \cdot (\mathfrak{p}) - \frac{\log(\mathcal{N}(\mathbf{a}))}{n} \sum_{\nu} (\nu). \quad (2.14)$$

Occasionally, we also use the non-normalized version of d^0 , called $d : \mathcal{I}_K^{\mathfrak{m}} \rightarrow \text{Div}_{K^{\mathfrak{m}}}$, which maps into $\text{Div}_{K^{\mathfrak{m}}}$ instead.

$$d : \mathcal{I}_K^{\mathfrak{m}} \rightarrow \text{Div}_{K^{\mathfrak{m}}}, \quad \mathbf{a} \mapsto \sum_{\mathfrak{p} \nmid \mathfrak{m}} \text{ord}_{\mathfrak{p}}(\mathbf{a}) \cdot (\mathfrak{p}).$$

The infinite part \mathbf{a}_{∞} of \mathbf{a} consists of a formal real sum of infinite places, which can be mapped into $K_{\mathbb{R}}$,

$$\text{Exp}(\cdot_{\infty}) : \text{Div}_{K^{\mathfrak{m}}}^0 \rightarrow K_{\mathbb{R}}, \quad \mathbf{a} \mapsto \text{Exp}(\mathbf{a}_{\infty}) = (e^{x_{\nu\sigma}})_{\sigma} \in K_{\mathbb{R}}.$$

2.4.3. Relation with Other Number-theoretic Groups

The groups and their relations treated above fit nicely in the diagram of exact sequences given in Figure 2.8, where the middle row sequence splits with the section d^0 . In this diagram we use the notations $\mathcal{O}_{K^{\mathfrak{m},1}}^{\times} = \mathcal{O}_K^{\times} \cap K^{\mathfrak{m},1}$, $\mu_{K^{\mathfrak{m},1}} = \mu_K \cap K^{\mathfrak{m},1}$ and $\text{Princ}_K^{\mathfrak{m}} = \{(\alpha) \mid \alpha \in K^{\mathfrak{m},1}\} \subseteq \mathcal{I}_K^{\mathfrak{m}}$. The group $\text{Cl}_K^{\mathfrak{m}}$ is called the ideal ray class group with respect to \mathfrak{m} and is defined by the exact sequence involved; the group $T^{\mathfrak{m}} = H/\Lambda_{K^{\mathfrak{m},1}}$ is the ‘logarithmic ray unit torus’, with $\Lambda_{K^{\mathfrak{m},1}} = \text{Log}(\mathcal{O}_{K^{\mathfrak{m},1}}^{\times}) = \{(\log |\sigma(\eta)|)_{\sigma} \mid \eta \in \mathcal{O}_{K^{\mathfrak{m},1}}^{\times}\}$.

$$\begin{array}{ccccccc}
& 0 & & 0 & & 0 & \\
\downarrow & & \downarrow & & \downarrow & & \\
0 \longrightarrow \mathcal{O}_{K^{\mathfrak{m},1}}^{\times}/\mu_{K^{\mathfrak{m},1}} \longrightarrow K^{\mathfrak{m},1}/\mu_{K^{\mathfrak{m},1}} \longrightarrow \text{Princ}_K^{\mathfrak{m}} \longrightarrow 0 & & & & & & \\
\text{Log} \downarrow & & (\cdot) \downarrow & & d^0 \downarrow & & \\
0 \longrightarrow H \longrightarrow \text{Div}_{K^{\mathfrak{m}}}^0 \xrightarrow{\text{a} \mapsto \text{Exp}(\mathbf{a}_f)} \mathcal{I}_K^{\mathfrak{m}} \longrightarrow 0 & & & & & & \\
\downarrow & & \downarrow & & \downarrow & & \\
0 \longrightarrow T^{\mathfrak{m}} \longrightarrow \text{Pic}_{K^{\mathfrak{m}}}^0 \longrightarrow \text{Cl}_K^{\mathfrak{m}} \longrightarrow 0 & & & & & & \\
\downarrow & & \downarrow & & \downarrow & & \\
0 & & 0 & & 0 & &
\end{array}$$

Figure 2.8.: A commutative diagram of short exact sequences involving the Arakelov ray class group.

The (ray) unit groups $\mathcal{O}_K, \mathcal{O}_{K^{\mathfrak{m},1}}^{\times}$, the (ray) class groups $\text{Cl}_K, \text{Cl}_K^{\mathfrak{m}}$, and the ray groups $K^{\mathfrak{m},1}$ and $K^{\mathfrak{m}}$ are tightly related by an exact sequence. With this exact sequence one can relate the (relative) cardinalities of these groups.

Lemma 2.15. *Let K be a number field and let $\mathfrak{m} \subseteq \mathcal{O}_K$ be any modulus. Then we have the following exact sequence of groups*

$$0 \rightarrow \mathcal{O}_{K^{\mathfrak{m},1}}^{\times} \rightarrow \mathcal{O}_K^{\times} \rightarrow K^{\mathfrak{m}}/K^{\mathfrak{m},1} \rightarrow \text{Cl}_K^{\mathfrak{m}} \rightarrow \text{Cl}_K \rightarrow 0.$$

In particular, $|\mathcal{O}_K^{\times}/\mathcal{O}_{K^{\mathfrak{m},1}}^{\times}| \cdot |\text{Cl}_K^{\mathfrak{m}}| = \phi(\mathfrak{m}) \cdot |\text{Cl}_K|$, where $\phi(\mathfrak{m}) = |K^{\mathfrak{m}}/K^{\mathfrak{m},1}| = |(\mathcal{O}_K/\mathfrak{m})^*|$.

Proof. By considering the kernel-cokernel exact sequence (see Figure A.1) of the commutative triangle

$$\begin{array}{ccc} & K^{\mathfrak{m}} & \\ \nearrow & & \searrow \\ K^{\mathfrak{m},1} & \xrightarrow{\quad} & \mathcal{I}_K^{\mathfrak{m}} \end{array}$$

one obtains the exact sequence

$$0 \rightarrow \mathcal{O}_{K^{\mathfrak{m},1}}^{\times} \rightarrow \mathcal{O}_K^{\times} \rightarrow K^{\mathfrak{m}}/K^{\mathfrak{m},1} \rightarrow \text{Cl}_K^{\mathfrak{m}} \rightarrow \text{Cl}_K \rightarrow 0,$$

where we use the fact that $\mathcal{I}_K^{\mathfrak{m}}/K^{\mathfrak{m}} \simeq \text{Cl}_K$ by the approximation theorem [Chi08, Ch. 3, Thm. 1.1]. In particular, one can ‘compress’ this sequence to an exact sequence of finite groups

$$0 \rightarrow \mathcal{O}_K^{\times}/\mathcal{O}_{K^{\mathfrak{m},1}}^{\times} \rightarrow K^{\mathfrak{m}}/K^{\mathfrak{m},1} \rightarrow \text{Cl}_K^{\mathfrak{m}} \rightarrow \text{Cl}_K \rightarrow 0,$$

yielding $|\mathcal{O}_K^{\times}/\mathcal{O}_{K^{\mathfrak{m},1}}^{\times}| \cdot |\text{Cl}_K^{\mathfrak{m}}| = |K^{\mathfrak{m}}/K^{\mathfrak{m},1}| \cdot |\text{Cl}_K|$. The isomorphism between $K^{\mathfrak{m}}/K^{\mathfrak{m},1}$ and $(\mathcal{O}_K/\mathfrak{m})^*$ follows from the following short exact sequence, where the map $K^{\mathfrak{m}} \rightarrow (\mathcal{O}_K/\mathfrak{m})^*$ sends $\kappa/\kappa' \in K^{\mathfrak{m}}$ to $(\kappa \bmod \mathfrak{m}) \cdot (\kappa' \bmod \mathfrak{m})^{-1} \in (\mathcal{O}_K/\mathfrak{m})^*$.

$$0 \rightarrow K^{\mathfrak{m},1} \rightarrow K^{\mathfrak{m}} \rightarrow (\mathcal{O}_K/\mathfrak{m})^* \rightarrow 0$$

□

One would expect that the ray unit torus $T^{\mathfrak{m}} = H/\text{Log}(\mathcal{O}_{K^{\mathfrak{m},1}}^{\times})$ and the unit torus $T = H/\text{Log}(\mathcal{O}_K^{\times})$ differ in volume by $|\mathcal{O}_K^{\times}/\mathcal{O}_{K^{\mathfrak{m},1}}^{\times}|$. This is true, up to a correction for whenever the modulus \mathfrak{m} causes $K^{\mathfrak{m},1}$ to have less roots of unity. This happens whenever $\zeta \not\equiv 1 \pmod{\mathfrak{m}}$ for some root of unity $\zeta \in K$.

Lemma 2.16. *Let K be a number field and let $H = \log K_{\mathbb{R}}^0$ be the hyperplane where the log unit lattice $\Lambda_K = \text{Log}(\mathcal{O}_K^{\times})$ and the log ray unit lattice $\Lambda_{K^{\mathfrak{m}}} = \text{Log}(\mathcal{O}_{K^{\mathfrak{m},1}}^{\times})$ live in. Then we have the following exact sequence*

$$0 \rightarrow \mu_{K^{\mathfrak{m},1}} \rightarrow \mu_K \rightarrow \mathcal{O}_K^{\times}/\mathcal{O}_{K^{\mathfrak{m},1}}^{\times} \rightarrow T^{\mathfrak{m}} \rightarrow T \rightarrow 0.$$

In particular, $|\mu_{K^{\mathfrak{m},1}}| \cdot |\mathcal{O}_K^{\times}/\mathcal{O}_{K^{\mathfrak{m},1}}^{\times}| = |\mu_K| \cdot \text{Vol}(T^{\mathfrak{m}})/\text{Vol}(T)$.

Proof. Applying the kernel-cokernel exact sequence to the following diagram yields the result.

$$\begin{array}{ccc}
& \mathcal{O}_K^\times & \\
\swarrow & & \searrow \\
\mathcal{O}_{K^{\mathfrak{m},1}}^\times & \xrightarrow{\quad} & H
\end{array}$$

□

2.4.4. The Volume of the Arakelov Ray Class Group

It will be proven useful to show that the volume of the Arakelov ray class group roughly follows the square root of the field discriminant times $\phi(\mathfrak{m}) = |(\mathcal{O}_K/\mathfrak{m})^*|$.

Lemma 2.17 (Volume of $\text{Pic}_{K^{\mathfrak{m}}}^0$). *For $n = [K : \mathbb{Q}] > 1$, we have*

$$\begin{aligned}
|\text{Pic}_{K^{\mathfrak{m}}}^0| &= |\text{Cl}_K^{\mathfrak{m}}| \cdot \text{Vol}(T^{\mathfrak{m}}) = \frac{|\mu_{K^{\mathfrak{m},1}}|}{|\mu_K|} \cdot \phi(\mathfrak{m}) \cdot h_K \cdot \text{Vol}(T) \\
&= \frac{|\mu_{K^{\mathfrak{m},1}}|}{|\mu_K|} \cdot \phi(\mathfrak{m}) h_K R_K \sqrt{n} 2^{-n\mathbb{C}/2},
\end{aligned} \tag{2.15}$$

and

$$\log |\text{Pic}_{K^{\mathfrak{m}}}^0| \leq \log \phi(\mathfrak{m}) + n \left(\frac{1}{2} \log(|\Delta_K|^{1/n}) + \log \log(|\Delta_K|^{1/n}) + 1 \right),$$

where $\phi(\mathfrak{m}) = |(\mathcal{O}_K/\mathfrak{m})^*|$. A simpler, derived bound is

$$\log(\text{Vol}(\text{Pic}_{K^{\mathfrak{m}}}^0)) \leq \log \mathcal{N}(\mathfrak{m}) + \log |\Delta_K|. \tag{2.16}$$

Proof. The first identity involving the volume of the Arakelov ray class group follows from the exact sequence in Figure 2.8. The second one can be deduced from the identities $|\text{Cl}_K^{\mathfrak{m}}| \cdot [\mathcal{O}_K : \mathcal{O}_{K^{\mathfrak{m},1}}^\times] = \phi(\mathfrak{m}) \cdot h_K$ and $\text{Vol}(T^{\mathfrak{m}}) = \text{Vol}(T) \cdot [\mathcal{O}_K^\times : \mathcal{O}_{K^{\mathfrak{m},1}}^\times] \cdot |\mu_{K^{\mathfrak{m},1}}| / |\mu_K|$ (see Lemmas 2.15 and 2.16). The third one follows from the volume computation of T in Lemma A.2.

The bound on the logarithm is obtained by using $\frac{|\mu_{K^{\mathfrak{m}},1}|}{|\mu_K|} \leq 1$, applying the class number formula [NS13, VII.§5, Cor. 5.11] and Louboutin's bound [Lou00] on the residue ρ_K of the Dedekind zeta function at $s = 1$:

$$\begin{aligned} |\mathrm{Pic}_{K^{\mathfrak{m}}}^0| &\leq \phi(\mathfrak{m})h_KR_K\sqrt{n}2^{-n_{\mathbb{C}}/2} = \frac{\phi(\mathfrak{m})\rho_K\sqrt{|\Delta_K|} \cdot |\mu_K| \cdot \sqrt{n}}{2^{n_{\mathbb{R}}}(2\sqrt{2\pi})^{n_{\mathbb{C}}}} \\ &\leq \phi(\mathfrak{m}) \cdot \sqrt{|\Delta_K|} \cdot \rho_K \leq \phi(\mathfrak{m})\sqrt{|\Delta_K|} \left(\frac{e \log |\Delta_K|}{2(n-1)}\right)^{n-1} \\ &\leq \phi(\mathfrak{m})\sqrt{|\Delta_K|} \left(\frac{e \log |\Delta_K|}{n}\right)^n, \end{aligned}$$

For the bound on the logarithm, we write

$$n \log(e \log |\Delta_K|/n) = n \log \log(|\Delta_K|^{1/n}) + n.$$

For the simpler bound in Equation (2.16) we use the fact that $\frac{e \log |x|}{|x|} \leq 1$ for all $x \in \mathbb{R}$. Therefore,

$$\frac{e \log \left(|\Delta_K|^{\frac{1}{2(n-1)}}\right)}{|\Delta_K|^{\frac{1}{2(n-1)}}} \leq 1,$$

and thus $\left(\frac{e \log |\Delta_K|}{2(n-1)}\right)^{n-1} \leq \sqrt{|\Delta_K|}$. □

We let $\mathcal{U}(\mathrm{Pic}_{K^{\mathfrak{m}}}^0) = \frac{1}{|\mathrm{Pic}_{K^{\mathfrak{m}}}^0|} \cdot \mathbf{1}_{\mathrm{Pic}_{K^{\mathfrak{m}}}^0}$ denote the uniform distribution over the Arakelov ray class group.

Fourier theory over the Arakelov ray class group

As the Arakelov ray class group $\mathrm{Pic}_{K^{\mathfrak{m}}}^0$ is a compact abelian group, every function in⁶ $L_2(\mathrm{Pic}_{K^{\mathfrak{m}}}^0) = \{f : \mathrm{Pic}_{K^{\mathfrak{m}}}^0 \rightarrow \mathbb{C} \mid \int_{\mathrm{Pic}_{K^{\mathfrak{m}}}^0} |f|^2 < \infty\}$ can be

⁶The measure on the Arakelov class group is unique up to scaling – it is the Haar measure. By fixing the volume of $\mathrm{Pic}_{K^{\mathfrak{m}}}^0$ as in Lemma 2.17, we fix this scaling as well. We use then *this* particular scaling of the Haar measure for the integrals over the Arakelov class group.

uniquely decomposed into a character sum

$$f = \sum_{\substack{\chi \in \widehat{\text{Pic}_{K^m}^0}}} a_\chi \cdot \chi,$$

with $a_\chi \in \mathbb{C}$. In the proof of Theorem 4.3, we will make use of Parseval's identity [DE16, Thm. 3.4.8] (see also Theorem 2.1) in the following form.

$$\int_{\text{Pic}_{K^m}^0} |f|^2 = \|f\|_2^2 = \frac{1}{|\text{Pic}_{K^m}^0|} \sum_{\chi \in \widehat{\text{Pic}_{K^m}^0}} |a_\chi|^2 \quad (2.17)$$

2.4.5. An Example of an Arakelov Class Group

We compute the Arakelov class group of a totally real cubic field. Let $K = \mathbb{Q}(\alpha)$ where $\alpha \in \mathbb{C}$ is defined by the polynomial

$$f(x) = x^3 - x^2 - 9x + 10. \quad (2.18)$$

Computing the ring of integers

The discriminant of this polynomial equals $\Delta(f) = 1957 = 19 \cdot 103 > 0$. Because this is square free, the ring of integers of K equals $\mathcal{O}_K = \mathbb{Z}[\alpha]$, and $\Delta_K = 1957$. Since the discriminant is positive, the cubic field must be totally real, by Brill's theorem. The Minkowski bound can then be computed as $M_K = \sqrt{|\Delta_K|} \cdot \frac{3!}{3^3} \approx 9.83$.

Computations in the class group

The class group is therefore generated by the primes with norm at most 9.83, which are the four prime ideals $\mathfrak{p}_2, \mathfrak{q}_2, \mathfrak{p}_5, \mathfrak{q}_5$. This can be seen by factoring the polynomial $f(x)$ modulo \mathbb{F}_p for $p = 2, 3, 5, 7$; noting that $f \pmod{3}$ and $f \pmod{7}$ are irreducible, and $f(x) \equiv x(x^2 + x + 1) \pmod{2}$ and $f(x) \equiv x(x^2 + 4x + 1) \pmod{5}$. We have $(2) = \mathfrak{p}_2\mathfrak{q}_2$ and $(5) = \mathfrak{p}_5\mathfrak{q}_5$, so, for the class group it is enough to consider only $\mathfrak{p}_2 = (2, \alpha)$ and $\mathfrak{p}_5 = (5, \alpha)$.

Additionally, we have $(\alpha) = \mathfrak{p}_2\mathfrak{p}_5$ and $(\alpha - 2) = \mathfrak{p}_2^2$. This can be seen by computing the norms of α and $\alpha - 2$, which equal $f(0) = 10$ and $f(2) = -4$ respectively. Since $(\alpha - 2) \subseteq (2, \alpha) = \mathfrak{p}_2$ we must have $(\alpha - 2) = \mathfrak{p}_2^2$. Combining these relations yields that the class group is generated by \mathfrak{p}_2 and is either trivial or of order 2. We will show that the latter is the case; for that we need the fundamental units.

Computing units and (a multiple of) the regulator

The elements $\alpha - 1$ and $\alpha - 3$ are units in \mathcal{O}_K , since $\mathcal{N}(\alpha - 1) = f(1) = 1$ and $\mathcal{N}(\alpha - 3) = f(3) = 1$. Under the Minkowski embedding, the element α sends to $(-3.04096, 1.12946, 2.9115)$, and 1 to $(1, 1, 1)$. Therefore, the images under the Minkowski embedding of $\alpha - 1$ and $\alpha - 3$ are respectively $\approx (-4.04096, 0.12946, 1.9115)$ and $\approx (-6.04096, -1.87054, -0.0885)$. Taking the Logarithmic image of the absolute values yields $\text{Log}(\alpha - 1) = (1.40, -2.04, 0.64)$ and $\text{Log}(\alpha - 3) = (1.80, 0.63, -2.42)$. Putting these vectors into a matrix, one obtains

$$B = \begin{bmatrix} 1.40 & -2.04 & 0.64 \\ 1.80 & 0.63 & -2.42 \end{bmatrix}, \quad (2.19)$$

of which the absolute determinant of any 2×2 minor equals 4.554, which must be an approximation of a multiple of the regulator R_K . So surely, $R_K \leq 4.554$.

Computing an approximation of the Dedekind residue

Computing an approximation of the residue of the Dedekind zeta function $\rho_K = \lim_{s \rightarrow 1} (s - 1)\zeta_K(s)$ by means of a truncated combined Euler product, we obtain

$$\rho_K \approx \frac{\prod_{p < 100} (1 - 1/p)}{\prod_{\mathcal{N}(\mathfrak{p}) < 100} (1 - 1/\mathcal{N}(\mathfrak{p}))} = 0.827.$$

By the class number formula (see Equation (2.11)), we have that

$$R_K h_K = \frac{\rho_K \cdot \sqrt{|\Delta_K|} \cdot |\mu_K|}{2^{n_{\mathbb{R}}} \cdot (2\pi)^{n_{\mathbb{C}}}} \approx \frac{0.827 \cdot 44.24 \cdot 2}{2^3 \cdot (2\pi)^0} = 9.15$$

Since $h_K \in \{1, 2\}$ and $R_K \leq 4.554$, we must have $h_K = 2$ and $R_K \approx 4.554$.

Assembling the Arakelov class group from the unit group and the class group

We have that $H = \{(x, y, z) \in \mathbb{R}^3 \mid x + y + z = 0\}$ equals the hyperplane where the logarithmic unit lattice lives in, and the log unit lattice equals $\Lambda_K = \mathcal{L}(B)$, where $\mathcal{L}(B)$ is the lattice generated by the rows of the matrix in Equation (2.19). The log unit torus is then equal to $T = H/\mathcal{L}(B)$.

The Arakelov class group Pic_K^0 of the cubic field K then has two connected components, one consisting of T , and one consisting of $T + [d^0(\mathfrak{p}_2)]$ (see Equation (2.14)). The maps of the exact sequence

$$0 \rightarrow T \rightarrow \text{Pic}_K^0 \rightarrow \text{Cl}_K \rightarrow 0$$

just consist of inclusion $T \hookrightarrow \text{Pic}_K^0$ and projection $\text{Pic}_K^0 \rightarrow \text{Cl}_K$, where $T \subseteq \text{Pic}_K^0$ sends to the trivial ideal class, and $T + [d^0(\mathfrak{p}_2)]$ sends to $[\mathfrak{p}_2] \in \text{Cl}_K$.

Computing elements in the Arakelov class group

We will compute the positions of $[d^0(\mathfrak{p}_5)]$, $[d^0(\mathfrak{q}_2)]$ and $[d^0(\mathfrak{p}_{17})]$ in the Arakelov class group, where $\mathfrak{p}_5 = (5, \alpha)$, $\mathfrak{q}_2 = (2, \alpha^2 + \alpha + 1)$ and $\mathfrak{p}_{17} = (17, \alpha + 1)$. This accounts to computing the discrete logarithm in the ideal class group and reducing modulo the logarithmic unit lattice.

As we have $\mathfrak{p}_5\mathfrak{p}_2 = (\alpha)$ and $\mathfrak{p}_2^2 = (\alpha - 2)$, we compute $\mathfrak{p}_5 = (\alpha)\mathfrak{p}_2^{-1} = (\alpha) \cdot (\alpha - 2)^{-1} \cdot \mathfrak{p}_2$. In terms of divisors, we have

$$\left\langle \frac{\alpha}{\alpha - 2} \right\rangle = \langle \mathfrak{p}_5 \rangle - \langle \mathfrak{p}_2 \rangle - \text{Log}(\alpha/(\alpha - 2)),$$

where we use the abbreviation $\text{Log}(\beta) = \sum_{\nu} \log |\sigma_{\nu}(\beta)| \cdot \langle \nu \rangle$. So,

$$\begin{aligned} d^0(\mathfrak{p}_5) &= \langle \mathfrak{p}_5 \rangle - \frac{1}{3} \cdot \text{Log}(5) = \langle \mathfrak{p}_2 \rangle + \left\langle \frac{\alpha}{\alpha - 2} \right\rangle + \text{Log}\left(\frac{\alpha}{\alpha - 2}\right) - \frac{1}{3} \cdot \text{Log}(5). \\ &= d^0(\mathfrak{p}_2) + \left\langle \frac{\alpha}{\alpha - 2} \right\rangle + \frac{1}{3} \cdot \text{Log}(2/5) + \text{Log}\left(\frac{\alpha}{\alpha - 2}\right). \end{aligned}$$

Taking Arakelov classes, thus letting vanish the part $(\alpha/(\alpha-2))$ (as it is a principal divisor), we obtain that

$$\begin{aligned} [d^0(\mathfrak{p}_5)] &= [d^0(\mathfrak{p}_2)] + \frac{1}{3} \cdot \text{Log}(2/5) + \text{Log}(\alpha/(\alpha-2)) \\ &\approx [d^0(\mathfrak{p}_2)] + (-0.81, -0.05, 0.86) \\ &\approx [d^0(\mathfrak{p}_2)] + (2.39, -1.46, -0.92) \in [d^0(\mathfrak{p}_2)] + T \end{aligned}$$

where the last computation just adds both rows of the logarithmic unit matrix from Equation (2.19) (in order to get in a fixed fundamental domain). A similar computation for \mathfrak{q}_2 , satisfying $\mathfrak{p}_2\mathfrak{q}_2 = (2)$, gives $(2/(\alpha-2)) = (\mathfrak{q}_2) - (\mathfrak{p}_2) + \text{Log}(2/(\alpha-2))$, and therefore

$$\begin{aligned} [d^0(\mathfrak{q}_2)] &= [d^0(\mathfrak{p}_2)] - \frac{1}{3} \text{Log}(2) + \text{Log}(2/(\alpha-2)) \\ &\approx [d^0(\mathfrak{p}_2)] + (-1.15, 0.60, 0.55) \\ &\approx [d^0(\mathfrak{p}_2)] + (2.05, -0.81, -1.23) \in [d^0(\mathfrak{p}_2)] + T. \end{aligned}$$

where, again, the last computation adds both rows of the logarithmic unit matrix from Equation (2.19). For $\mathfrak{p}_{17} = (17, \alpha+1)$, compute the norm of $\alpha+1$ to see that it equals 17, therefore, $(\alpha+1) = (\mathfrak{p}_{17}) - \text{Log}(\alpha+1)$. This implies

$$\begin{aligned} [d^0(\mathfrak{p}_{17})] &= -\frac{1}{3} \text{Log}(17) + \text{Log}(\alpha+1) \approx (-0.23, -0.19, 0.42) \\ &\approx (1.57, 0.44, -2.00) \in T \end{aligned}$$

where the last computation adds the last row of the logarithmic unit matrix from Equation (2.19).

The Arakelov classes of the primes \mathfrak{q}_2 , \mathfrak{p}_5 and \mathfrak{p}_{17} are portrayed in Figures 2.9 and 2.10, in which the full Arakelov class group of $K = \mathbb{Q}(\alpha)$ is displayed. In Figure 2.9, the primes are visualized in a two-dimensional fundamental domain (a disjoint union of two parallelograms) whereas in Figure 2.10 the toroidal nature of the Arakelov class group is exemplified.

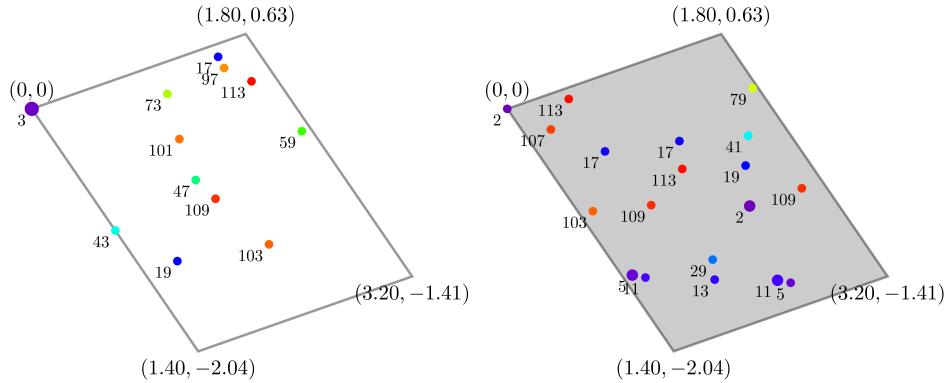


Figure 2.9.: In this picture, the Arakelov class group of $K = \mathbb{Q}(\alpha)$ is portrayed, where $\alpha \in \mathbb{C}$ is defined by the polynomial $f(x) = x^3 - x^2 - 9x + 10$. Due to the fact that the class group has order 2 and the unit group is free of rank 2, the Arakelov class group can be portrayed as a disjoint union of two parallelograms, serving as a fundamental domain. The connected component of the unit $[\mathcal{O}_K]$ is the white parallelogram on the left-hand side; the gray parallelogram is associated with the non-trivial ideal class group element. Prime ideals up to norm 113 are displayed as points, where the color hue varies with the size of the associated prime number, and the size of the point with the residue class degree of the prime ideal. The prime ideal $q_2 = (2, \alpha^2 + \alpha + 1)$ of residue class degree 2 can be seen in the gray parallelogram as the rather large dot labeled with '2'. The prime ideal $p_5 = (5, \alpha)$ is located at the right bottom of the gray parallelogram, as a purple point. The prime ideal $p_{17} = (17, \alpha + 1)$ is principal and it is therefore located in the white parallelogram, at the top right corner, as a blue point.

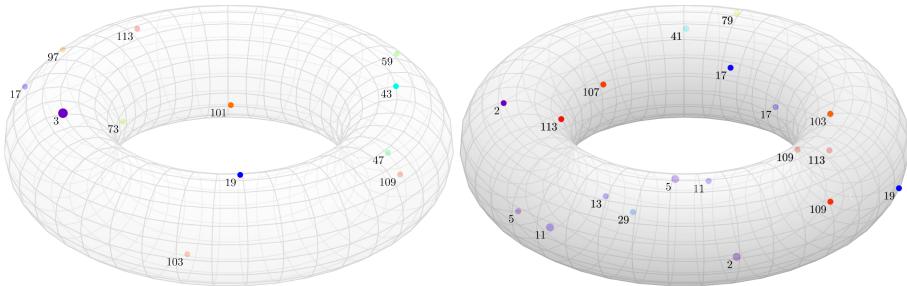


Figure 2.10.: This picture shows the Arakelov class group of the same number field K as in Figure 2.9. One obtains this image by ‘gluing’ the gray parallelogram into a gray torus and the white parallelogram into a white torus from Figure 2.9. The prime ideals with norms up to 113 are displayed accordingly. Note that the location of the smaller prime ideals seem to be skewed on the gray torus; but as the norms increase, the division among the two tori, but also on the tori seem to get more and more uniform. This phenomenon can be seen as a manifestation of the random walk theorem, which states that from a certain lower bound on the norms, prime ideals become more and more uniformly located on these tori; assuming the extended Riemann hypothesis (see Theorem 4.3).

2.5. Lattices

2.5.1. General Lattices

A lattice Λ is a discrete subgroup of a real vector space. In the following, we assume that this real vector space has dimension m and that the lattice is full-rank, i.e., $\text{span}(\Lambda)$ equals the whole real space. A lattice can be represented by a basis $B = (\mathbf{b}_1, \dots, \mathbf{b}_m)$ such that $\Lambda = \{\sum_i x_i \mathbf{b}_i, x_i \in \mathbb{Z}\}$. Important notions in lattice theory are the covolume $\text{Vol}(\Lambda)$, which equals the volume of the hypertorus $\text{span}(\Lambda)/\Lambda$ (alternatively, $\text{Vol}(\Lambda)$ is the absolute determinant of any basis of Λ); the first minimum $\lambda_1(\Lambda) = \min_{v \in \Lambda \setminus \{0\}} \|v\|$; and the last minimum $\lambda_m(\Lambda)$, which equals the minimal radius $r > 0$ such that $\{v \in \Lambda \mid \|v\| \leq r\}$ is of full rank m . The equivalent notions with respect to the maximum norm $\|\cdot\|_\infty$ instead of the Euclidean norm are denoted by $\lambda_1^{(\infty)}(\Lambda)$ and $\lambda_m^{(\infty)}(\Lambda)$. We will also use the following notation for the covering radius; $\text{cov}_2(\Lambda)$ (and $\text{cov}_\infty(\Lambda)$ for the maximum norm analogue), which is the minimum $r > 0$ such that any element $x \in \text{span}(\Lambda)$ is at most r -close to a lattice point.

For any (full-rank) lattice $\Lambda \subseteq \mathbb{R}^m$ we denote by $\Lambda^* = \{v \in \mathbb{R}^m \mid \langle v, \ell \rangle \in \mathbb{Z} \text{ for all } \ell \in \Lambda\}$ the *dual lattice* of Λ . It is a lattice of full rank and, furthermore, for any basis B of Λ holds that $D = (B^T)^{-1}$ is a basis of Λ^* .

We will be interested into the following algorithmic problem over lattices.

Definition 2.18 (γ -Hermite Shortest Vector Problem). *Given as input a basis of a rank m lattice Λ , the problem γ -Hermite-SVP consists in computing a non-zero vector $v \in \Lambda$ such that*

$$\|v\| \leq \gamma \cdot \text{Vol}(\Lambda)^{1/m}.$$

2.5.2. Divisors and Ideal Lattices

It will be proven useful to view both ideals and Arakelov divisors as lattices in the real vector space $K_{\mathbb{R}}$, where $K_{\mathbb{R}}$ has its (Euclidean or maximum)

norm inherited from the complex vector space it lives in. Explicitly, the Euclidean and maximum norm of $\alpha \in K$ are respectively defined by the rules $\|\alpha\|_2^2 = \sum_{\sigma} |\sigma(\alpha)|^2$ and $\|\alpha\|_{\infty} = \max_{\sigma} |\sigma(\alpha)|$, where σ ranges over all embeddings $K \rightarrow \mathbb{C}$. By default, $\|\alpha\|$ refers to the Euclidean norm $\|\alpha\|_2$.

For any ideal \mathfrak{a} of K , we define the associated lattice $\mathfrak{a} \subseteq K_{\mathbb{R}}$ to be the image of $\mathfrak{a} \subseteq K$ under the Minkowski embedding, which is clearly a discrete subgroup of $K_{\mathbb{R}}$. By slightly abusing the notation we both denote the ideal and the associated lattice with the same symbol \mathfrak{a} . In particular, \mathcal{O}_K is a lattice and we will always assume throughout this thesis (except stated otherwise) that we know a \mathbb{Z} -basis $(\mathbf{b}_1, \dots, \mathbf{b}_n)$ of \mathcal{O}_K . For Arakelov divisors $\mathbf{a} = \sum_{\mathfrak{p}} n_{\mathfrak{p}} \cdot (\mathfrak{p}) + \sum_{\nu} x_{\nu} \cdot (\nu)$, the associated lattice is defined as follows.

$$\text{Exp}(\mathbf{a}) = \left\{ (e^{x_{\nu\sigma}} \cdot \sigma(\alpha))_{\sigma} \mid \alpha \in \prod \mathfrak{p}^{n_{\mathfrak{p}}} \right\} = \text{diag}((e^{x_{\nu\sigma}})_{\sigma}) \cdot \prod \mathfrak{p}^{n_{\mathfrak{p}}} \subseteq K_{\mathbb{R}},$$

where diag denotes a diagonal matrix. We have $\text{Vol}(\mathfrak{a}) = \sqrt{|\Delta_K|} \mathcal{N}(\mathfrak{a})$ for ideals $\mathfrak{a} \in \mathcal{I}_K$ and, for Arakelov divisors $\mathbf{a} \in \text{Div}_K$,

$$\text{Vol}(\text{Exp}(\mathbf{a})) = \sqrt{|\Delta_K|} \cdot \prod_{\sigma} e^{x_{\nu\sigma}} \cdot \mathcal{N}(\prod_{\mathfrak{p}} \mathfrak{p}^{n_{\mathfrak{p}}}) = \sqrt{|\Delta_K|} \cdot e^{\deg(\mathbf{a})}.$$

The associated lattice $\text{Exp}(\mathbf{a})$ of a divisor is of a special kind, which we call *ideal lattices*, as in the following definition.

Definition 2.19 (Ideal lattices). *Let K be a number field with ring of integers \mathcal{O}_K . An ideal lattice of K is a \mathcal{O}_K -module $I \subseteq K_{\mathbb{R}}$, with the additional requirement that there exists an $x \in K_{\mathbb{R}} \setminus \{0\}$ such that $xI \subseteq \mathcal{O}_K$. We denote the group of ideal lattices by IdLat_K .*

Note that the lattices \mathfrak{a} for $\mathfrak{a} \in \mathcal{I}_K$ are special cases of ideal lattices, which we will call *fractional ideal lattices*. Since the Minkowski embedding is injective, the Minkowski embedding provides a bijection between the set of fractional ideals and the set of fractional ideal lattices.

The set IdLat_K of ideal lattices forms a group; the product of two ideal lattices $I = x\mathfrak{a}$ and $J = y\mathfrak{b}$ is defined by the rule $I \cdot J = xy\mathfrak{a}\mathfrak{b}$. It is clear that $\mathcal{O}_K \subseteq K_{\mathbb{R}}$ is the unit ideal lattice and $x^{-1}\mathfrak{a}^{-1}$ is the inverse ideal lattice of $x\mathfrak{a}$.

The map $\text{Exp}(\cdot) : \text{Div}_K^0 \rightarrow \text{IdLat}_K$, $\mathbf{a} \mapsto \text{Exp}(\mathbf{a})$ sends an Arakelov divisor to an ideal lattice. The image under this map is the following subgroup of IdLat_K .

$$\text{IdLat}_K^0 = \{x\mathbf{a} \mid \mathcal{N}(\mathbf{a}) \prod_{\sigma} x_{\sigma} = 1 \text{ and } x_{\sigma} > 0 \text{ for all } \sigma\}.$$

Definition 2.20 (Isometry of ideal lattices). *For two ideal lattices $L, L' \in \text{IdLat}_K^0$, we say that L and L' are K -isometric, denoted by $L \sim L'$, when there exists $(\xi_{\sigma}) \in K_{\mathbb{R}}$ with $|\xi_{\sigma}| = 1$ such that $(\xi_{\sigma})_{\sigma} \cdot L = L'$.*

In other words, we consider two ideal lattices to be K -isometric if they only differ in component-wise complex phase. Being K -isometric is an equivalence relation on IdLat_K^0 that is compatible with the group operation.

Relation between ideal lattices and Arakelov classes

Denoting Iso_K for the subgroup $\{L \in \text{IdLat}_K^0 \mid L \sim \mathcal{O}_K\} \subset \text{IdLat}_K^0$, we have the following result.

Lemma 2.21 (Arakelov classes are ideal lattices up to isometry). *Denoting $P : \text{IdLat}_K^0 \rightarrow \text{Pic}_K^0$ for the map $x\mathbf{a} \longmapsto \sum_{\mathfrak{p}} \text{ord}_{\mathfrak{p}}(\mathbf{a})[\mathfrak{p}] + \sum_{\nu} \log(x_{\sigma_{\nu}})[\nu]$ modulo principal divisors, we have the following exact sequence.*

$$0 \rightarrow \text{Iso}_K \rightarrow \text{IdLat}_K^0 \xrightarrow{P} \text{Pic}_K^0 \rightarrow 0.$$

Proof. This is a well-known fact (e.g., [Sch08]), but we give a proof for completeness. It suffices to show that P is a well-defined surjective homomorphism and its kernel is Iso_K . In order to be well-defined, P must satisfy $P(x\mathbf{a}) = P(x'\mathbf{a}')$ whenever $x\mathbf{a} = x'\mathbf{a}'$. Assuming the latter, we obtain $x^{-1}x'\mathcal{O}_K = (\mathbf{a}')^{-1}\mathbf{a} = \alpha\mathcal{O}_K$, for some $\alpha \in K^*$, as the module is a free \mathcal{O}_K -module. This implies that $(x^{-1}x')_{\sigma} = \sigma(\eta\alpha)$ for all embeddings $\sigma : K \rightarrow \mathbb{C}$, for some unit $\eta \in \mathcal{O}_K^{\times}$. Therefore, we have, $P(x\mathbf{a}) - P(x'\mathbf{a}') = \sum_{\mathfrak{p}} \text{ord}_{\mathfrak{p}}(\alpha)[\mathfrak{p}] + \sum_{\nu} \log((x_{\sigma_{\nu}})^{-1}x'_{\sigma_{\nu}})[\nu] = \langle \eta\alpha \rangle$; i.e., their difference is a principal divisor, meaning that their image in Pic_K^0 is the same.

One can check that P is a homomorphism, and its surjectivity can be proven by constructing an ideal lattice in the pre-image of a representative divisor $\mathbf{a} = \sum_{\mathfrak{p}} n_{\mathfrak{p}}[\mathfrak{p}] + \sum_{\nu} x_{\nu}[\nu] \in \text{Div}_K^0$ of an Arakelov class $[\mathbf{a}]$, e.g., $(e^{x_{\nu\sigma}})_{\sigma} \cdot \prod_{\mathfrak{p}} \mathfrak{p}^{n_{\mathfrak{p}}}$.

We finish the proof by showing that the kernel of P indeed equals Iso_K . Suppose $x\mathfrak{a} \in \ker(P)$, i.e., $P(x\mathfrak{a}) = \sum_{\mathfrak{p}} \text{ord}_{\mathfrak{p}}(\mathfrak{a})[\mathfrak{p}] + \sum_{\nu} \log(x_{\sigma\nu})[\nu] = \langle \alpha \rangle$ is a principal divisor. This means that $\mathfrak{a} = \alpha \mathcal{O}_K$ and $x = (|\sigma(\alpha)|^{-1})_{\sigma}$, i.e., $x\mathfrak{a} = (|\sigma(\alpha)|^{-1})_{\sigma} \alpha \mathcal{O}_K = \left(\frac{\sigma(\alpha)}{|\sigma(\alpha)|}\right)_{\sigma} \cdot \mathcal{O}_K$, so $x\mathfrak{a} \sim \mathcal{O}_K$, implying $x\mathfrak{a} \in \text{Iso}_K$. This shows that $\ker P \subseteq \text{Iso}_K$. The reverse inclusion starts with the observation that $x\mathfrak{a} \sim \mathcal{O}_K$ directly implies that $\mathfrak{a} = \alpha \mathcal{O}_K$ is principal, by the fact that $x\mathfrak{a}$ is a free \mathcal{O}_K -module. So, $(x_{\sigma}\sigma(\alpha))_{\sigma} \cdot \mathcal{O}_K = x\alpha \mathcal{O}_K = (\xi_{\sigma})_{\sigma} \cdot \mathcal{O}_K$ for some $(\xi_{\sigma})_{\sigma} \in K_{\mathbb{R}}$ with $|\xi_{\sigma}| = 1$. Therefore, $|x_{\sigma}\sigma(\eta\alpha)| = |\xi_{\sigma}| = 1$, i.e., $|x_{\sigma}| = |\sigma(\eta\alpha)|^{-1}$ for some unit $\eta \in \mathcal{O}_K^{\times}$. From here one can directly conclude that $P(x\mathfrak{a}) = P((|\sigma(\eta\alpha)|^{-1})_{\sigma} \alpha \mathcal{O}_K) = \langle \eta\alpha \rangle$, a principal divisor. \square

Bounds on invariants of ideal lattices

Denote $\Gamma(\Lambda) = \lambda_n(\Lambda)/\lambda_1(\Lambda)$, and define, for a fixed number field K :

$$\Gamma_K = \sup_{\mathbf{a} \in \text{Div}_K} \Gamma(\text{Exp}(\mathbf{a})) \quad (2.20)$$

Recall the notion of the covering radius; $\text{cov}_2(\Lambda)$ (and $\text{cov}_{\infty}(\Lambda)$ for the maximum norm), which is the minimum $r > 0$ such that any element $x \in \text{span}(\Lambda)$ is at most r -close to a lattice point. For ideal lattices, we then do have the following useful bounds, which are used often throughout this thesis.

Lemma 2.22. *For any modulus $\mathfrak{m} \subseteq \mathcal{O}_K$ and any divisor $\mathbf{a} \in \text{Div}_{K^{\mathfrak{m}}}$,*

- (i) $\Gamma_K \leq \lambda_n^{(\infty)}(\mathcal{O}_K) \leq |\Delta_K|^{1/n}$;
- (ii) *For cyclotomic number fields K , $\Gamma_K = 1$;*
- (iii) $\lambda_n(\text{Exp}(\mathbf{a})) \leq \sqrt{n} \cdot \Gamma_K \cdot \text{Vol}(\text{Exp}(\mathbf{a}))^{1/n}$;
- (iv) $\text{cov}_{\infty}(\text{Exp}(\mathbf{a})) \leq \text{cov}_2(\text{Exp}(\mathbf{a})) \leq n/2 \cdot \Gamma_K \cdot \text{Vol}(\text{Exp}(\mathbf{a}))^{1/n}$.

Proof. The bound $\lambda_n^{(\infty)}(\mathcal{O}_K) \leq |\Delta_K|^{1/n}$ can be proven by means of the techniques of [Bha+20, Thm. 3.1], as is done in Theorem A.4 of Appendix A.1. To obtain the bound $\Gamma_K \leq \lambda_n^{(\infty)}(\mathcal{O}_K)$, pick an arbitrary divisor $\mathbf{a} \in \text{Div}_{K^m}$ and choose a shortest element $x\alpha \in \text{Exp}(\mathbf{a})$ with $x = \text{Exp}(\mathbf{a}_\infty)$ and $\alpha \in \text{Exp}(\mathbf{a}_f) \in \mathcal{I}_K^m$. That means $\|x\alpha\| = \lambda_1(\text{Exp}(\mathbf{a}))$. Then $\text{Exp}(\mathbf{a}) \supset x \cdot (\alpha)$, and therefore

$$\lambda_n(\text{Exp}(\mathbf{a})) \leq \lambda_n(x \cdot (\alpha)) \leq \|x\alpha\|_2 \cdot \lambda_n^{(\infty)}(\mathcal{O}_K) = \lambda_1(\text{Exp}(\mathbf{a})) \cdot \lambda_n^{(\infty)}(\mathcal{O}_K),$$

which proves part (i). Part (ii) follows from part (i) and the fact that $\|\zeta\| = \|1\|$ for roots of unity $\zeta \in K$. Part (iii) is essentially Minkowski's bound $\lambda_1(\text{Exp}(\mathbf{a})) \leq \sqrt{n} \text{Vol}(\text{Exp}(\mathbf{a}))^{1/n}$ combined with the definition of Γ_K . The last item follows from the fact that $\text{cov}_2(\Lambda) \leq \sqrt{n}/2 \cdot \lambda_n(\Lambda)$ [Mic]. \square

2.5.3. The Gaussian Function and Smoothing Errors

Let n be a fixed positive integer. For any parameter $s > 0$, we consider the n -dimensional *Gaussian function*

$$\rho_s^{(n)} : \mathbb{R}^n \rightarrow \mathbb{C}, \quad x \mapsto e^{-\frac{\pi\|x\|^2}{s^2}},$$

(where we drop the (n) whenever it is clear from the context), which is well known to have the following basic properties.

Lemma 2.23. *For all $s > 0$, $n \in \mathbb{N}$ and $x, y \in \mathbb{R}^n$, we have $\int_{z \in \mathbb{R}^n} \rho_s(z) dz = s^n$, $\mathcal{F}_{\mathbb{R}^n}\{\rho_s\} = \int_{y \in \mathbb{R}^n} \rho_s(y) e^{-2\pi i \langle y, \cdot \rangle} dy = s^n \rho_{1/s}$, $\rho_s(x)^2 = \rho_{s/\sqrt{2}}(x)$. and $\sqrt{\rho_s(x)\rho_s(y)} = \rho_{2s}(x+y)\rho_{2s}(x-y)$.*

Remark 2.24. *From these properties it follows that the the L_2 -norm of $x \mapsto s^{m/2} \cdot \sqrt{\rho_{1/s}(x)}$ equals 1, i.e., $\|s^{m/2} \cdot \sqrt{\rho_{1/s}(x)}\|_{\mathbb{R}^m}^2 = 1$.*

The following two results (and the variations we discuss below) will play an important role and will be used several times in this paper: *Banaszczyk's*

bound, originating from [Ban93], and the *smoothing parameter*, as introduced by Micciancio and Regev [MR07]. They allow us to control

$$\rho_s(X) := \sum_{x \in X} \rho_s(x),$$

for certain discrete subsets $X \subseteq \mathbb{R}^m$. For ease of notation, we let

$$\beta_z^{(n)} := \left(\frac{2\pi e z^2}{n} \right)^{n/2} e^{-\pi z^2},$$

which decays super-exponentially in z (for fixed n). In particular, we have $\beta_t^{(n)} \leq e^{-t^2}$ for all $t \geq \sqrt{n}$. The following formulation of Banaszczyk's lemma is obtained from [MS18, Eq. (1.1)].

Lemma 2.25 (Banaszczyk's Bound). *Whenever $r/s \geq \sqrt{\frac{n}{2\pi}}$,*

$$\rho_s((\Lambda + t) \setminus \mathcal{B}_r) \leq \beta_{r/s}^{(n)} \cdot \rho_s(\Lambda),$$

where $\mathcal{B}_r = \mathcal{B}_r(0) = \{x \in \mathbb{R}^n \mid \|x\|_2 < r\}$.

Definition 2.26 (Smoothing parameter). *Given an $\varepsilon > 0$ and a lattice Λ , the smoothing parameter $\eta_\varepsilon(\Lambda)$ is the smallest real number $s > 0$ such that $\rho_{1/s}(\Lambda^* \setminus \{0\}) \leq \varepsilon$. Here, Λ^* is the dual lattice of Λ .*

Lemma 2.27 (Smoothing Error). *Let $\Lambda \in \mathbb{R}^n$ be a full rank lattice, and let $s \geq \eta_\varepsilon(\Lambda)$. Then, for any $t \in \mathbb{R}^n$,*

$$(1 - \varepsilon) \frac{s^n}{\det \Lambda} \leq \rho_s(\Lambda + t) \leq (1 + \varepsilon) \frac{s^n}{\det \Lambda}. \quad (2.21)$$

We have the following two useful upper bounds for full-rank n -dimensional lattices Λ [MR07, Lm. 3.2 and 3.3]: $\eta_\varepsilon(\Lambda) \leq \sqrt{\log(2n(1 + 1/\varepsilon))} \cdot \lambda_n(\Lambda)$ for all $\varepsilon > 0$ and $\eta_1(\Lambda) \leq \eta_{2^{-n}}(\Lambda) \leq \sqrt{n}/\lambda_1(\Lambda^*) \leq \sqrt{n} \cdot \lambda_n(\Lambda)$. The latter leads to the following corollary.

Corollary 2.28. *Let L be an ideal lattice in IdLat_K . Let $t \in \mathbb{R}^n$ be arbitrary and $s \geq n \cdot \lambda_n(\mathcal{O}_K) \cdot \text{Vol}(L)^{1/n}$. Then it holds that*

$$\left| \frac{\rho_s(L - t) \cdot \text{Vol}(L)}{s^n} - 1 \right| \leq 2^{-n}, \quad (2.22)$$

Proof. By the assumption on s and by Lemma 2.22, we have $s \geq n \cdot \lambda_n(\mathcal{O}_K) \cdot \text{Vol}(L)^{1/n} \geq \sqrt{n} \cdot \lambda_n(L) \geq \eta_{2^{-n}}(\Lambda)$. The result follows then from Lemma 2.31. \square

Alternative descriptions of the smoothing bound

Imitating techniques from Micciancio and Regev [MR07, Lm. 3.2], we have:

Lemma 2.29. *Let $s \geq \frac{\sqrt{m}}{\lambda_1(\Lambda^*)}$. Then $\rho_{1/s}(\Lambda^* \setminus \{0\}) \leq 2 \cdot \beta_{s\lambda_1(\Lambda^*)}$.*

As a direct corollary, we have the following result.

Corollary 2.30. *Let $s \geq 2\sqrt{m}$, and let $x \in \mathbb{R}^m$ with $\|x\|_\infty \leq 1/2$. Then*

$$\rho_{1/s}(\mathbb{Z}^m \setminus \{0\} + x) \leq 2 \cdot \beta_{s/2}.$$

Proof. We have $\rho_{1/s}(\mathbb{Z}^m \setminus \{0\} + x) \leq \rho_{1/s}((\mathbb{Z}^m + x) \setminus \mathcal{B}_{\frac{1}{2}}) \leq \beta_{s/2} \cdot \rho_{1/s}(\mathbb{Z}^m)$, where the second inequality follows from Lemma 2.25. Using Lemma 2.29 to argue that $\rho_{1/s}(\mathbb{Z}^m) = 1 + \rho_{1/s}(\mathbb{Z}^m \setminus \{0\}) \leq 1 + 2 \cdot \beta_s \leq 2$ then proves the claim. \square

The following lemma, which combines [MR07, Lm. 4.1] and [MR07, Lm. 3.2], controls the fluctuation of the sum $\rho_s(\Lambda + t)$ for varying $t \in \mathbb{R}^m$.

Lemma 2.31 (Smoothing Error). *Let $\Lambda \in \mathbb{R}^m$ be a full rank lattice, and let $s \geq \sqrt{m}/\lambda_1(\Lambda^*)$. Then, for any $t \in \mathbb{R}^m$,*

$$(1 - 2 \cdot \beta_{s\lambda_1(\Lambda^*)}) \frac{s^m}{\det \Lambda} \leq \rho_s(\Lambda + t) \leq (1 + 2 \cdot \beta_{s\lambda_1(\Lambda^*)}) \frac{s^m}{\det \Lambda}. \quad (2.23)$$

Corollary 2.32. For $s \geq \frac{\sqrt{m}}{\lambda_1(\Lambda^*)}$ and for any $t \in \mathbb{R}^m$, we have $\rho_s(\Lambda + t) \leq 2 \frac{s^m}{\det \Lambda}$.

Proof. Using Lemma 2.31 and noticing $2 \cdot \beta_{s\lambda_1(\Lambda^*)} \leq 2 \cdot \beta_{\sqrt{m}} \leq 1$ yields the result. \square

2.5.4. Gaussian Distributions

In this work, both discrete and continuous Gaussian distributions play a major role. We denote both of these distributions with $\mathcal{G}_{X,s,c}$, where the subscript X is a metric space which supports the distribution and thus indicates whether the Gaussian is discrete or continuous. More concretely, for discrete spaces X like lattices, $\mathcal{G}_{X,s}$ a discrete Gaussian, whereas for continuous spaces it is a continuous Gaussian. For the cases of a vector space and a lattice, the definition is spelled out below.

Continuous Gaussian distribution. For a real vector space H of dimension n , a parameter $s > 0$ and a center $c \in H$, we write $\mathcal{G}_{H,s,c}$ the continuous Gaussian distribution over H with density function $\rho_s(x - c)/s^n$ for all $x \in H$. When the center c is 0, we simplify the notation as $\mathcal{G}_{H,s}$.

Discrete Gaussian distributions. For any lattice $L \subset \mathbb{R}^n$, we define the discrete Gaussian distribution over L of standard deviation $s > 0$ and center $c \in \mathbb{R}^n$ by

$$\forall x \in L, \mathcal{G}_{L,s,c} = \frac{\rho_s(x - c)}{\rho_s(L - c)}.$$

When the center c is 0, we simplify the notation as $\mathcal{G}_{L,s}$.

2.6. The Lipschitz Condition

Theorem 2.33 (Rademacher's theorem). A Lipschitz function $\mathbf{f} : \mathbb{R}^m/\Lambda \rightarrow \mathcal{H}$ has weak partial derivatives $\partial_{x_j} \mathbf{f} : \mathbb{R}^m/\Lambda \rightarrow \mathcal{H}$ lying in $L_2(\mathbb{R}^m/\Lambda)$. In

particular,

$$\sum_{j=1}^m \|\partial_{x_j} \mathbf{f}\|_{\mathbb{R}^m/\Lambda}^2 \leq \text{Lip}(\mathbf{f})^2.$$

Proof. Combining the proof of [Hei04, Thm. 4.1 and 4.9] and [Vil85, Thm. 2] on measures of compact sets, we obtain this result. \square

Corollary 2.34. *Let $\mathbf{f} : \mathbb{R}^m/\Lambda \rightarrow \mathcal{H}$ be a Lipschitz-continuous function, and denote by $|c_{\ell^*}\rangle$ the vectorial Fourier coefficients of \mathbf{f} . Then,*

$$\sum_{\substack{\ell^* \in \Lambda^* \\ \|\ell^*\| \geq B}} \langle c_{\ell^*} | c_{\ell^*} \rangle \leq \frac{\text{Lip}(\mathbf{f})^2}{4\pi^2 B^2}.$$

Proof. Since \mathbf{f} is Lipschitz, we can apply Theorem 2.33. Furthermore, the identity $|\mathbf{f}(x)\rangle = \sum_{\ell^* \in \Lambda^*} |c_{\ell^*}\rangle e^{2\pi i \langle x, \ell^* \rangle}$ implies that

$$|\partial_{x_j} \mathbf{f}(x)\rangle = 2\pi i \sum_{\ell^* \in \Lambda^*} \ell_j^* |c_{\ell^*}\rangle e^{2\pi i \langle x, \ell^* \rangle}$$

almost everywhere ([Wer07, Lm. V.2.11] or [RA08, Lm. 2.16]). Finally, given that $\sum_{j=1}^m \|\partial_{x_j} \mathbf{f}\|_{\mathbb{R}^m/\Lambda}^2 \leq \text{Lip}(\mathbf{f})^2$, Plancherel's identity implies that

$$\begin{aligned} \text{Lip}(\mathbf{f})^2 &\geq \sum_{j=1}^m \|\partial_{x_j} \mathbf{f}\|_{\mathbb{R}^m/\Lambda}^2 = 4\pi^2 \sum_{\ell^* \in \Lambda^*} \|\ell^*\|_2^2 \cdot \langle c_{\ell^*} | c_{\ell^*} \rangle \\ &\geq 4\pi^2 \sum_{\substack{\ell^* \in \Lambda^* \\ \|\ell^*\|_2 \geq B}} \|\ell^*\|_2^2 \cdot \langle c_{\ell^*} | c_{\ell^*} \rangle \geq 4B^2 \pi^2 \sum_{\substack{\ell^* \in \Lambda^* \\ \|\ell^*\|_2 \geq B}} \langle c_{\ell^*} | c_{\ell^*} \rangle, \end{aligned}$$

from which the claim follows. \square

3. The Continuous Hidden Subgroup Problem

3.1. Summary

This chapter is about a complexity analysis of a slightly modified algorithm of Eisenträger et al. [Eis+14] that quantumly solves the *continuous hidden subgroup problem*. This problem consists of finding a ‘hidden lattice’ Λ in \mathbb{R}^m given a (possibly) quantum function $\mathbf{f} : \mathbb{R}^m \rightarrow \mathcal{S}$ that is periodic with respect to this lattice Λ . This computational problem falls into the class of the so-called ‘period finding problems’.

This quantum algorithm mimics the blueprint of Shor’s algorithm for finding a hidden subgroup H in a discrete abelian group G , given an oracle function on the group that is strictly periodic with respect to H . This blueprint consists of consecutively sampling a uniform quantum superposition over all group elements, applying an oracle call to the H -periodic function, and computing a discrete quantum Fourier transform. Then, one measures to obtain a character $\chi \in \hat{G}$ that has H in its kernel. Assembling enough of such characters allows to retrieve H itself.

The quantum algorithm solving the continuous hidden subgroup problem

The quantum algorithm of this chapter deviates from this blueprint in a few ways. (1) Since the ambient group \mathbb{R}^m is continuous, we need to cut-off and discretize this space to get something finite and thus processable by a quantum computer. This has as a consequence that the Fourier transform

becomes discretized as well, inducing errors with respect to the continuous Fourier transform. (2) The initial state of this quantum algorithm does not consist of a uniform quantum sample but of a Gaussian state instead. This is done to ease the analysis, as both the Gaussian function and its Fourier transform (which is also a Gaussian function) have tight tail bounds. (3) The measurement output is, due to the cut-off and discretization, always an *approximation* of a dual lattice vector $\ell^* \in \Lambda^*$ (which can be seen as a character of \mathbb{R}^m with Λ in its kernel). So, in the end, we cannot expect more to gain from this algorithm than an *approximate basis* \tilde{B} of the lattice Λ . (4) Such an approximate basis is obtained as follows. By sampling many approximate $\ell^* \in \Lambda^*$, LLL-reducing these samples to an approximate basis \tilde{D} of the dual lattice Λ^* , and inverting and transposing \tilde{D} , one retrieves an approximate basis \tilde{B} of Λ .

Analysis of the algorithm

Each deviation from the original ‘hidden subgroup problem blueprint’ causes difficulties; mostly those difficulties take the shape of *discretization errors*. We show how to solve these difficulties per deviation. Tackling these difficulties was already partially done by Eisenträger et al. [Eis+14]; we revisit their work to obtain a more explicit and precise complexity.

- (1) The discrete Fourier transform and the continuous (real) Fourier transform can be shown to differ not too much if their input function is *continuous enough*. A large part of this chapter (Section 3.5) is devoted to show that if the Λ -periodic oracle function is *Lipschitz continuous*, the induced error by using a discrete Fourier transform instead of a continuous one can be reasonably bounded.
- (2) For the initial input to be Gaussian, one needs to know how to actually assemble this state on a quantum computer. Such a Gaussian superposition has already been shown to be computable in polynomial time by Kitaev and Webb [KW08], but for completeness we included a more precise complexity estimate in Appendix A.5.

(3) Due to the discrete nature of the quantum algorithm, the output dual lattice point can only be approximated within a certain distance. The maximum allowed distance (relative to the minimum distance $\lambda_1(\Lambda^*)$ of the dual lattice Λ^*) will be a parameter in the algorithm, called $\delta > 0$.

One of the problems that might occur is that the output dual lattice points are not equidistributed enough on Λ^* , thus not giving enough information to retrieve a basis of Λ . An extra assumption on the Λ -periodic function \mathbf{f} is needed to avoid such a situation; which we call *separating*. A separating Λ -periodic function can be intuitively thought of as being not too constant. Showing that such an oracle will yield equidistributed points in Λ^* is the object of Section 3.6.

(4) From many such δ -close dual lattice points one can compute an approximate basis of the dual lattice Λ^* by means of LLL-reduction; from this approximate dual basis one can obtain a basis of Λ by inversion and transposition. These operations (LLL-reduction and inversion) are quite *numerical unstable*, meaning that they make existing errors in the input progressively larger. Using a result of Buchmann and Kessler [BK96] one can reasonably bound the final error (see Section 3.7).

Relation with the Arakelov (ray) class group

The computation of the Arakelov (ray) class group can be phrased in terms of a hidden lattice problem; a fact that can already be inferred from the original applications of the hidden lattice problem, namely computing (S)-unit groups and class groups [BS16; Eis+14] in works of Biasse, Song and Eisenträger et al. By a slight modification in formulation of the ideas in these papers one can construct an oracle on the Arakelov divisor group that is periodic with respect to the principal divisors. In this modification, a ‘reduced’ version of the Arakelov (ray) divisor group is used, one with only finitely many prime ideals, that are required to generate the ideal class group. Finding the periodicity of this oracle then allows to find explicit relations that define the Arakelov (ray) class group.

3. The Continuous Hidden Subgroup Problem

At the time of writing, a precise complexity estimation (beyond polynomial time) of the oracle function in this approach to quantumly compute Arakelov (ray) class groups is still open.

3.2. Introduction

The Hidden Subgroup Problem

Among all quantum algorithms, Shor’s algorithm [Sho94] for factoring and finding discrete logarithms is singular by its cryptanalytic implications. Due to progress toward the realization of large quantum computers, this celebrated algorithm is now motivating the standardization of quantum-resistant schemes [Nat17], in preparation of a global update of widely deployed encryption and authentication protocols.

The core idea of quantum period finding [Sho94] is not limited to factoring and discrete logarithm. The *Hidden Subgroup Problem*, formalized in [ME98], serves as a convenient interface between the quantum-algorithmic techniques for period finding, and applications to solve other computational problems, in particular problems arising from number theory. We will here discuss only the case of commutative groups. The cases of non-abelian groups such as dihedral groups are very interesting as well and have fascinating connections with lattice problems [Reg04b]; however, no polynomial time algorithm is known for those cases, and the best known algorithm has sub-exponential complexity [Kup05], using very different techniques.

The simplest version of the Hidden Subgroup Problem consists of finding a hidden subgroup H in a *finite* abelian group G , when given access to a strictly H -periodic function $\mathbf{f} : G \rightarrow S$. Here, in the language of representation theory, the off-the-shelf period-finding quantum algorithm finds a uniformly random character $\chi \in \hat{G}$ that acts trivially on H . Shor’s original algorithm [Sho94] for integer factoring finds a hidden subgroup H in the ambient group \mathbb{Z} . The infiniteness of \mathbb{Z} induces some “cut-off” error;

nevertheless, the distribution of the algorithm’s output is still concentrated around the multiples of the inverse period.

A generalization to the real line $H = \mathbb{R}$ was given by Hallgren [Hal07] and allows to solve Pell’s equation. The case of real vector space of constant dimension $H = \mathbb{R}^c$ has also been studied [Hal05; SV05], and permits the computation of unit groups of number fields of fixed finite degree.

The *Continuous* Hidden Subgroup Problem

The latest generalization of the HSP algorithm, given by Eisenträger, Hallgren, Kitaev and Song in an extended abstract [Eis+14], targets the ambient group $G = \mathbb{R}^m$ (for a non-constant dimension m) with a hidden discrete subgroup $H = \Lambda$, i.e. a *lattice*. Next to the ambient group \mathbb{R}^m being *continuous*, an additional special feature is that the Λ -periodic function \mathbf{f} is assumed to produce a “quantum output”. More formally, $\mathbf{f} : \mathbb{R}^m \rightarrow \mathcal{S}$, $x \mapsto |\mathbf{f}(x)\rangle$, where \mathcal{S} is the state space of a quantum system, and the HSP algorithm is given access to a unitary that maps $|x\rangle|0\rangle$ to $|x\rangle|\mathbf{f}(x)\rangle$. A crucial observation here is that $|\mathbf{f}(x)\rangle$ and $|\mathbf{f}(y)\rangle$ are *not* necessarily orthogonal (or even distinct) for distinct x and y modulo Λ . In other words, it is not assumed that \mathbf{f} is *strictly* periodic, but merely that $|\mathbf{f}(x)\rangle$ and $|\mathbf{f}(y)\rangle$ are “somewhat orthogonal” for x and y that are “not too close” modulo Λ , and that \mathbf{f} is Lipschitz continuous.

More specifically, they consider a variation of the standard HSP algorithm in order to tackle the Continuous Hidden Subgroup Problem (CHSP). In order to deal with the continuous nature of the domain \mathbb{R}^m of \mathbf{f} , the given HSP algorithm acts on a bounded “grid” of points within \mathbb{R}^m . Additionally, the algorithm is modified in the following ways: (1) The initial state is not a uniform superposition (over the considered grid points in \mathbb{R}^n) but follows a trigonometric distribution, and (2) the quantum Fourier transform is done “remotely”, i.e., rather than applying it to the actual register, the register is entangled with an ancilla and the quantum Fourier transform is then applied to the ancilla instead. According to Eisenträger et al. [Eis+14], applying the

3. The Continuous Hidden Subgroup Problem

quantum Fourier transform directly would make the resulting approximation errors difficult to analyze.

As an application, Eisenträger et al. also gave a quantum polynomial time algorithm for computing the unit group of a number field in their article [Eis+14]. This was generalized by Biasse and Song [BS16] to the computation of S -unit groups, and therefore to the computation of class groups and to finding a generator of a principal ideals. This led to solving the shortest vector problem in certain ideal lattices for non-trivial approximation factors [Cra+16; CDW17; PHS19]. While the cryptanalytic consequences for ideal-lattice based cryptography seem limited so far [DPW19], these results demonstrate a hardness gap between ideal lattices and general ones.

Our Contributions

The goal of this chapter is to provide a complete, modular, and quantitative analysis of (a slightly modified version of) the Continuous HSP quantum algorithm given by [Eis+14]. More concretely, we provide an explicit bound on the number of qubits needed by the algorithm, clarifying the dependency on the parameters of the Continuous HSP instance and on the required precision and success probability. This shows explicitly in what parameters the algorithm is polynomial time and with what exponent.

The algorithm that we consider and analyze differs from the one considered in [Eis+14] in the following points:

- First, we specify the initial state of the algorithm to have Gaussian amplitudes, while [Eis+14, Sec. 6.2] suggests to use a cropped trigonometric function; as far as we can see, our choice makes the analysis simpler and tighter thanks to the well known tail-cut and smoothness bounds of Banaszczuk [Ban93] and Micciancio and Regev [MR07].
- Secondly, we do not make use of a “remote” Fourier transform but instead follow the blueprint of Shor’s original algorithm in that respect; the claimed advantage of the “remote” Fourier transform is unclear to us.

These modifications simplify the algorithm and its analysis. Due to the lack of details given in [Eis+14], we can not state a complexity comparison, but we think this variation is at least as efficient as the original algorithm.

Our analysis is divided into four parts, each summarized by a formal statement given in Sections 3.3.3 to 3.3.6, leading to the main theorem (Section 3.3.2). We insist on this modular presentation, so as to enable future work on optimization and specialization of this algorithm to instances of interests; specific suggestions follow.

Dual lattice sampling. In the first part, which is the technically more involved one, we show that the appropriately discretized and finitized, but otherwise (almost) standard HSP quantum algorithm produces sample points in \mathbb{R}^m that lie close to the dual lattice Λ^* with high probability. More precisely, and more technically speaking, we show that the algorithm’s output is a sample point close to $\ell^* \in \Lambda^*$ with probability close to $\langle c_{\ell^*} | c_{\ell^*} \rangle$, where the vectors $|c_{\ell^*}\rangle$ are the Fourier coefficients of the function \mathbf{f} . This is in line with the general HSP approach, where for instance Shor’s algorithm for period finding over \mathbb{Z} produces a point that is close to a random multiple of the inverse period, except with bounded probability.

In this first part (Section 3.4 and Section 3.5), we bound the complexity of the core algorithm in terms of the error probability that we allow in the above context of a sampling algorithm, and depending on the Lipschitz constant of \mathbf{f} . In particular, we show that the number of qubits grows as mQ , where Q , the “number of qubits per dimension”, grows linearly in the logarithm of the Lipschitz constant of \mathbf{f} , the logarithm of the inverse of the error probability and the logarithm of the inverse of the (absolute) precision, and quasi-linearly in m . The running time of the algorithm is then bounded by $O(mQ \log(mQ))$, by using an approximate Fourier transform [HH00].

Full dual recovery. In the second part, Section 3.6, we then relate the parameters of the Continuous HSP instance to the number of sample points,

3. The Continuous Hidden Subgroup Problem

and thus to how often the core algorithm needs to be repeated, necessary in order to have an approximation of the entire dual lattice Λ^* .

Primal basis reconstruction. In the third part, Section 3.7, we study the numerical stability of reconstructing an approximate basis of the primal lattice Λ from a set of approximate generators of the dual lattice Λ^* . This is based on the Buchmann-Pohst algorithm [BK96] already mentioned in [Eis+14]. The claim of [Eis+14] involves intricate quantities related to sublattices of Λ , making the final complexity hard to derive; we provide a simpler statement with a detailed proof.

Gaussian state preparation. Finally, in Appendix A.5, we revisit the quantum polynomial-time algorithm for the preparation of the Gaussian initial state [GR02; KW08] used as a black-box in our first part, and provide its precise complexity.

Conclusion. These four parts lead to our formal and quantitative version of the informal CHSP Theorem of Eisenträger et al. [Eis+14, Thm. 6.1], stated as Theorem 3.3 in Section 3.3.2.

Conclusion and Research Directions

Our conclusion is that, in its generic form, the Continuous Hidden Subgroup Problem is rather expensive to solve; not accounting for other parameters than the dimension m , it already requires $\tilde{O}(m^3)$ qubits and $\tilde{O}(m^4)$ quantum gates (using an approximate quantum Fourier transform). However, this inefficiency seems to be a consequence of its genericity. In particular, the core algorithm for Dual Lattice Sampling would only need $\tilde{O}(m^2)$ qubits, if it wasn't for accommodating for the terrible numerical stability of the Primal Basis Reconstruction step. Similarly, we expect the number of samples needed to generate the dual lattice to be significantly smaller for smoother oracle functions.

All in all, our modular analysis of the generic steps of the CHSP algorithm sets the stage for analyzing and optimizing its specializations, in particular to cryptanalytic applications [Cra+16; CDW17]. We propose a few research directions towards this objective:

- Study the costs (qubits, quantum gates) and the parameters of the oracle functions from [Eis+14; BS16; Son13] for solving the Unit Group Problem, the Principal Ideal Problem (PIP), and for the computation of the class group.
- Find stronger hypotheses satisfied by the above oracle functions (or by variant thereof) that improve this generic analysis of the CHSP algorithm; or resort to an ad-hoc analysis of the Full Dual Recovery step by directly studying the spectrum of these oracle functions.
- Explore the possibility of a trade-off between the (classical) Primal Basis Reconstruction step and the (quantum) Dual Lattice Sampling step, possibly up to small sub-exponential classical complexity. More specifically, does replacing LLL by BKZ with a medium block-size substantially improve the numerical stability of Buchmann-Pohst algorithm?
- Exploit prior knowledge of sublattices (potentially close to full-rank) of the hidden lattice to accelerate or skip the Full Dual Recovery and Primal Basis Reconstruction steps. This is for example the case when solving PIP [BS16] while already knowing the unit group and the class group of a given number field. This would be applicable in the context of [Cra+16; CDW17].
- Exploit known symmetries of the hidden sublattice to improve the Full Dual Recovery and Primal Basis Reconstruction steps. Such symmetries are for example induced by the Galois action on the log-unit lattice and the lattice of class relation, in particular in the case of the cyclotomic number fields. This would again be applicable in the context of [Cra+16; CDW17].

Remark 3.1. *Recovering the exact hidden lattice is outside the scope of this work, since this task is application-dependent. It is even true that one*

cannot generally expect the quantum algorithm of this chapter to recover the exact hidden lattice, without extra information about this hidden lattice.

For instance, when applying this algorithm to compute the unit group \mathcal{O}_K^\times of a number field K , the hidden lattice will be the so-called logarithmic unit lattice. Of this lattice it is known that any point is of the shape $\text{Log}(\eta) = (\log |\sigma(\eta)|)_\sigma \in \text{Log } K_{\mathbb{R}}^0$ with $\eta \in \mathcal{O}_K^\times \subseteq \mathcal{O}_K$; its entries are logarithms of integral elements in a given number field. This is the extra information that is to be exploited in order to get the exact lattice. Namely, from a sufficiently good approximation of the logarithm of a unit one can obtain the exact underlying unit, simply by taking the exponential and rounding it to the closest element in the ring of integers \mathcal{O}_K .

3.3. Problem Statements and Results

3.3.1. Notation and Set-up

Here and throughout this chapter, \mathcal{H} is a complex Hilbert space of dimension $N = 2^n$, and \mathcal{S} is the unit sphere in \mathcal{H} ; thus, a vector in \mathcal{S} describes the state of a system of n qubits. For an arbitrary positive integer m , we consider a function

$$\mathbf{f} : \mathbb{R}^m \rightarrow \mathcal{S} \subset \mathcal{H}, x \mapsto |\mathbf{f}(x)\rangle$$

that is periodic with respect to a full rank lattice $\Lambda \subset \mathbb{R}^m$; hence, \mathbf{f} may be understood as a function $\mathbb{R}^m/\Lambda \rightarrow \mathcal{S}$. The function \mathbf{f} is assumed to be Lipschitz continuous with Lipschitz constant

$$\text{Lip}(\mathbf{f}) = \inf \left\{ L > 0 \mid \|\langle \mathbf{f}(x) \rangle - \langle \mathbf{f}(y) \rangle\|_{\mathcal{H}} \leq L \|x - y\|_{2, \mathbb{T}^m} \right\}.$$

Later, we will also require \mathbf{f} to be “sufficiently non-constant”. One should think of \mathbf{f} as an oracle that maps a classical input x to a quantum state over n qubits, which is denoted $|\mathbf{f}(x)\rangle$.

We write Λ^* for the dual lattice of Λ . By $\lambda_1(\Lambda)$ we denote the length of a shortest non-zero vector of Λ , and correspondingly for $\lambda_1(\Lambda^*)$. Since Λ is

typically clear from the context, we may just write λ_1 and λ_1^* instead of $\lambda_1(\Lambda)$ and $\lambda_1(\Lambda^*)$.

We denote by $\mathcal{B}_r(x) = \{y \in \mathbb{R}^m \mid \|y - x\| < r\}$ the open Euclidean ball with radius r around x . For the open ball around 0 we just denote \mathcal{B}_r , and for a set $X \subset \mathbb{R}^m$ we write $\mathcal{B}_r(X) = \bigcup_{x \in X} \mathcal{B}_r(x)$.

Definition 3.2 (Definition 1.1 from [Eis+14]). *A function $\mathbf{f} : \mathbb{R}^m \rightarrow \mathcal{S} \subset \mathcal{H}$ is said to be an (a, r, ϵ) -HSP oracle of the full-rank lattice $\Lambda \subset \mathbb{R}^m$ if*

- \mathbf{f} is Λ -periodic,
- \mathbf{f} is a -Lipschitz: $\text{Lip}(f) \leq a$,
- \mathbf{f} is (r, ϵ) -separating (see Figure 3.1): I.e., $|\langle \mathbf{f}(x) | \mathbf{f}(y) \rangle| \leq \epsilon$ for all $x, y \in \mathbb{R}^m$ satisfying $d_{\mathbb{R}^m / \Lambda}(x, y) \geq r$.

where $d_{\mathbb{R}^m / \Lambda}(x, y) = \min_{v \in \Lambda} \|x - y - v\|$ denotes the distance induced by the Euclidean distance of \mathbb{R}^n modulo Λ .

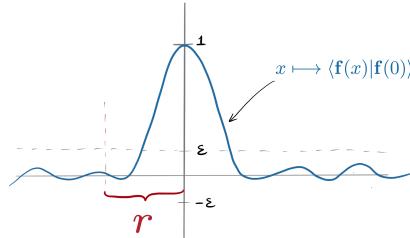


Figure 3.1.: A picture of what an (r, ϵ) -separating function \mathbf{f} should look like: outside of the interval of length $2r$ around the origin, the inner product $x \mapsto \langle \mathbf{f}(x) | \mathbf{f}(0) \rangle$ deviates from 0 by no more than ϵ .

3.3.2. Main Theorem: Continuous Hidden Subgroup Problem

Theorem 3.3. *There exists a quantum algorithm that, given access to an (a, r, ϵ) -HSP oracle with period lattice Λ , $r < \lambda_1(\Lambda)/6$ and $\epsilon < 1/4$, computes, with constant success probability, an approximate basis $\tilde{B} = B + \Delta_B$ of this lattice Λ , satisfying $\|\Delta_B\| < \tau$.*

3. The Continuous Hidden Subgroup Problem

This algorithm makes k quantum oracle calls to the (a, r, ϵ) -HSP oracle, and uses $mQ + n$ qubits, $O(kmQ \cdot (\log(kmQ))^2)$ quantum gates and $\text{poly}(m, \log \frac{a}{\lambda_1^*}, \log \frac{a}{\tau})$ classical bit operations, where

$$Q = O(mk) + O\left(\log \frac{a}{\lambda_1^*}\right) + O\left(\log \frac{1}{\lambda_1^* \cdot \tau}\right), \quad (3.24)$$

$$k = O\left(m \cdot \log\left(\sqrt{m} \cdot a \cdot (\det \Lambda)^{1/m}\right)\right) \quad (3.25)$$

Remark 3.4. Note that the quantities inside logarithms are homogeneous. In particular, scaling the lattice Λ by a factor f , also scales τ , $1/a$ and $1/\lambda_1^*$ by the same factor f , leaving the complexity parameters Q and k unaffected.

Remark 3.5. The expert reader may expect the “distortion” parameter $\lambda_1 \cdot \lambda_1^*$ of the lattice Λ to have a bearing on the complexity of this algorithm. It is indeed implicitly the case: the assumption the HSP definition implies that $ar \geq \text{Lip}(\mathbf{f}) \cdot r \geq 1 - \epsilon$ (see Figure 3.2), and therefore the theorem’s hypothesis requires $a \geq \text{Lip}(\mathbf{f}) \geq \frac{9}{4\lambda_1}$.

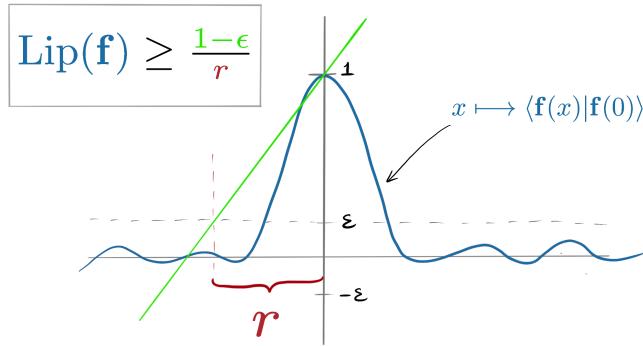


Figure 3.2.: Due to the (r, ϵ) -separating property of the oracle function \mathbf{f} , its Lipschitz constant cannot be too small.

The proof of Theorem 3.3 can be found in Section 3.3.7.

3.3.3. Dual Lattice Sampling Problem

Following our modular approach as outlined in the introduction, we first consider the following *Dual Lattice Sampling Problem*. Informally, the task is to sample points in \mathbb{R}^m that are respectively close to points $\ell^* \in \Lambda^*$ that follow the distribution $\mathcal{D}_{ideal}(\ell^*) = \langle c_{\ell^*} | c_{\ell^*} \rangle$, where $|c_{\ell^*}\rangle$ are the vectorial Fourier coefficients of $\mathbf{f} : \mathbb{R}^m / \Lambda \rightarrow \mathcal{S}$ (see Section 2.2.4).

Problem 3.6 (Dual Lattice Sampling Problem). *Given error parameter $\eta > 0$ and a relative distance parameter $\frac{1}{2} > \delta > 0$, and given oracle access to an HSP oracle \mathbf{f} as above, sample according to a (finite) distribution \mathcal{D} on \mathbb{R}^m that satisfies, for any $S \subseteq \Lambda^*$,*

$$p_S := \mathcal{D}(\mathcal{B}_{\delta\lambda_1^*}(S)) \geq \left(\sum_{\ell^* \in S} \langle c_{\ell^*} | c_{\ell^*} \rangle \right) - \eta. \quad (3.26)$$

In the problem statement above, $\mathcal{D}(\mathcal{B}_{\delta\lambda_1^*}(S))$ denotes the cumulative weight of the set $\mathcal{B}_{\delta\lambda_1^*}(S) = \bigcup_{s \in S} \mathcal{B}_{\delta\lambda_1^*}(s)$ with respect to the distribution \mathcal{D} . Here, $\mathcal{B}_{\delta\lambda_1^*}(s) = \{y \in \mathbb{R}^m \mid \|s - y\| < \delta\lambda_1^*\}$ is the open ball of radius $\delta\lambda_1^*$ around $s \in S \subseteq \Lambda^* \subseteq \mathbb{R}^m$.

Theorem 3.7. *Algorithm 2 solves the Dual Lattice Sampling Problem with parameters η and δ ; it uses one call to the Gaussian superposition subroutine (see Theorem 3.12), one quantum oracle call to \mathbf{f} , $mQ + n$ qubits, and $O(mQ \log(mQ))$ quantum gates, where*

$$Q = O(m \log(m)) + O\left(\log\left(\frac{a}{\eta \cdot \delta\lambda_1^*}\right)\right). \quad (3.27)$$

Remark 3.8. *Note that this step only requires smoothness of the HSP oracle (via the Lipschitz constant), but does not rely on the “separateness” assumption (third item of Definition 3.2). Indeed this third assumption will only play a role to ensure that those samples are actually non-trivial and usable.*

3.3.4. Full Dual Lattice Recovery

Recovering the full lattice (or, equivalently, its dual) requires an extra assumption on the oracle function \mathbf{f} , as captured by the third condition in the following definition, reformatted from Definition 1.1 of [Eis+14].

According to Eisenträger et al. [Eis+14], for (some undetermined) adequate parameters, Definition 3.2 ensures that the distribution on the dual lattice Λ^* is not concentrated on any proper sublattice, hence sufficiently many samples will generate the lattice fully. We formalize and quantify this proof strategy, and obtain the following quantitative conclusion. We note that the constraints on r and ϵ are milder than one could think, for example ϵ does not need to tend to 0 as a function of n or m . More precisely, a constant $\epsilon < 1/4$ and a constant $r \leq \lambda_1(\Lambda)/6$ would suffice.

Theorem 3.9. *Let $\mathbf{f} : \mathbb{R}^m \rightarrow \mathcal{S}$ be an (a, r, ϵ) -HSP oracle of the full-rank lattice $\Lambda \subset \mathbb{R}^m$, with $r \leq \lambda_1(\Lambda)/6$ and $\epsilon < 1/4$. Let $\mathcal{D}_\mathbf{f}$ be the distribution supported by Λ^* , with weight $\langle c_{\ell^*} | c_{\ell^*} \rangle$ at $\ell^* \in \Lambda^*$, where $|c_{\ell^*}\rangle$ are the vectorial Fourier coefficients of the function \mathbf{f} .*

Then, with overwhelming probability, we need at most

$$O\left(m \log_2 (ma \cdot \det(\Lambda)^{1/m})\right)$$

samples from $\mathcal{D}_\mathbf{f}$ to fully generate the lattice Λ^ .*

The above theorem is obtained by combining Lemma 3.21 and proposition 3.24 from Section 3.6, instantiating the parameter R to $R^2 = ma^2$. This choice is somewhat arbitrary and given for concreteness, however it does not have a critical quantitative impact.

3.3.5. Primal Basis Reconstruction

Theorem 3.10. *There exists a polynomial time algorithm, that, for any matrix $G \in \mathbb{R}^{k \times m}$ of k generators of a (dual) lattice Λ^* , and given an*

approximation $\tilde{G} = G + \Delta_G \in \mathbb{Q}^{k \times n}$, computes an approximation $\tilde{B} = B + \Delta_B$ of a basis B of the primal lattice Λ , such that

$$\|\Delta_B\|_\infty \leq \frac{2^{O(mk)} \cdot \|\tilde{G}\|_\infty^{m+1}}{(\lambda_1^*)^3 \cdot \det(\Lambda^*)} \cdot \|\Delta_G\|_\infty,$$

under the assumption that $\|\Delta_G\|_\infty < \frac{\min(1, (\lambda_1^*)^2) \cdot \det(\Lambda^*)}{2^{O(km)} \cdot \|\tilde{G}\|_\infty^{m+1}}$.

Remark 3.11. More specifically, the algorithm from Theorem 3.10 essentially consists of the Buchmann-Pohst algorithm [BP89; BK96] and a matrix inversion. Its complexity is dominated by two calls to LLL on matrices of dimension $(m+k) \times k$ and entry bit size $O(k^2 \log(\|\tilde{G}\|/\lambda_1^*))$ (see the discussion before [BK96, Cor. 4.1]). One can optimize the final running time by choosing a fast variant of LLL, e.g., [NS16].

Our contribution on this step is merely a completed numerical analysis, with the help of a theorem from [CSV12]. A claim with a similar purpose is given in [Eis+14], yet involves more intricate lattice quantities.

3.3.6. Gaussian State Preparation

The main algorithm of this paper requires the preparation of a multidimensional Gaussian initial state, which can be obtained by generating the one-dimensional Gaussian state on m parallel quantum registers. This task is known to be polynomial time [GR02; KW08], and we provide a quantitative analysis in Appendix A.5. The precise running time of preparing this Gaussian state is summarized below.

Theorem 3.12. For $q = 2^Q \in \mathbb{N}$, error parameter $\eta \in (0, 1)$ and $s > 2\sqrt{\log(m/\eta)}$, there exists a quantum algorithm that prepares the higher-dimensional Gaussian state

$$\frac{1}{\sqrt{\rho_{1/s}(\mathbb{D}_{\text{rep}}^m)}} \sum_{\mathbf{x} \in \mathbb{D}_{\text{rep}}^m} \sqrt{\rho_{1/s}(\mathbf{x})} |\mathbf{x}\rangle = \bigotimes_{j=1}^m \frac{1}{\sqrt{\rho_{1/s}(\frac{1}{q}[q]_c)}} \sum_{x \in \frac{1}{q}[q]_c} \sqrt{\rho_{1/s}(x)} |x\rangle,$$

within trace distance η , using $O(mQ + \log(\eta^{-1}))$ qubits and using $O(mQ \cdot \log(mQ\eta^{-1})^2)$ quantum gates.

Remark 3.13. In Theorem 3.3, we chose η to be $1/k^2$. Therefore, one call to the m -dimensional Gaussian state preparation with the parameters of Theorem 3.3 takes $O(mQ + \log(k))$ qubits and $O(mQ \log(kmQ)^2)$ quantum gates. As Theorem 3.3 requires k subsequent preparations of the m -dimensional Gaussian state, the total costs of the Gaussian state preparation steps are $O(mQ + \log(k))$ qubits (by reusing qubits) and $O(kmQ \log(kmQ)^2)$ quantum gates.

This is slightly more than the costs of k times applying the Fourier transform, and it explains the quantum gate complexity of $O(kmQ \log(kmQ)^2)$ in Theorem 3.3.

3.3.7. Proof of the Main Theorem

Proof of Theorem 3.3. The result is obtained by running Algorithm 1 and instantiating Theorems 3.7, 3.9, 3.10 and 3.12.

Correctness of Algorithm 1. In step one, the dual sampling algorithm (Algorithm 2) is applied k times with error probability $\eta = 1/k^2$. The probability that all measurements are actually $\delta\lambda_1^*$ -close to dual lattice points and are of length less than $\sqrt{m}a$ is then at least $(1 - \eta)^k = (1 - 1/k^2)^k \geq 1 - 1/k$, which is at least a constant success probability. We assume in the rest of the proof that all measurements are indeed $\delta\lambda_1^*$ -close to dual lattice points and of length less than $\sqrt{m} \cdot a$.

In step two, these $\delta\lambda_1^*$ -close-to- Λ^* samples are assembled into a matrix $k \times m$ -matrix \tilde{G} , on which is then applied the Buchmann-Pohst algorithm [BK96; BP89] twice. Subsequently, the resulting square matrix is inverted and transposed. By Theorem 3.10, this procedure runs in polynomial time and has no error probability. Due to the choice of δ and the fact that $\|\tilde{G}\|_\infty \leq \sqrt{m}a$ and $\|\tilde{G} - G\| < \delta \cdot \lambda_1^*$, we can apply Theorem 3.10 to obtain $\|\Delta_B\|_\infty = \|B - \tilde{B}\| < \tau$, as required. Note that the size of δ is chosen in such a way that the decline in precision (see Theorem 3.10) is taken care of. By Theorem 3.3, the matrix \tilde{G} indeed approximates a full generating set of Λ^* with overwhelming probability; implying that the output matrix \tilde{B}

Algorithm 1: Quantum algorithm that solves the continuous hidden subgroup problem

Require:

- An (a, r, ϵ) -oracle $\mathbf{f} : \mathbb{R}^m \rightarrow \mathcal{H}$ that is periodic with respect to the full-rank hidden lattice $\Lambda \subseteq \mathbb{R}^m$, whose dual lattice Λ^* has first minimum $\lambda_1^* = \lambda_1(\Lambda^*)$. We require the parameters ϵ and r to satisfy $\epsilon < 1/4$ and $r \leq \lambda_1(\Lambda)/6$.
- An error parameter τ quantifying the maximum allowed deviation of the output basis \tilde{B} from an actual basis B of Λ .

Ensure: With constant probability, an τ -approximated basis \tilde{B} of the lattice Λ . In other words, a matrix $\tilde{B} \in \text{Mat}_{m \times m}(\mathbb{Q})$ satisfying $\|\tilde{B} - B\| < \tau$ for some basis $B \in \text{Mat}_{m \times m}(\mathbb{R})$ of Λ , i.e., τ -close in the maximum norm induced matrix norm.

- 1: Apply the **dual sampling algorithm** (Algorithm 2) k times, with failure probability $\eta = 1/k^2$, Gaussian deviation $s = O(\sqrt{m \log(\eta^{-1})})$ and $V = O(\frac{m \log(\eta^{-1})}{\delta \lambda_1^*})$, where $k = O(m \log[\sqrt{m} \cdot a \cdot \det(\Lambda)^{1/m}])$ and $\delta = 2^{-O(mk)} \cdot (\sqrt{m} \cdot a)^{-(m+1)} \cdot \det(\Lambda)^{-1} \cdot (\lambda_1^*)^2 \cdot \tau$.
- 2: Assemble the k samples from above algorithm into a matrix \tilde{G} , **apply the Buchmann-Pohst algorithm twice** (see Section 3.7), and invert and transpose the resulting basis, yielding a matrix \tilde{B} .
- 3: **return** \tilde{B} .

approximates a basis of Λ with overwhelming probability (and not a basis of a strict sublattice of Λ).

Complexity estimate. We focus first on the less important complexity, the classical complexity. This complexity is mainly driven by LLL-algorithm and inversion in step (2) of Algorithm 1. This complexity can be bounded polynomially in the dimensions and the entry sizes of the matrix involved. The dimensions of \tilde{G} are $k \times m$, and can therefore be polynomially bounded by m , $\log a$ and $\log(\det \Lambda)$. The entry sizes (taking a precision of at least δ into account) can be polynomially bounded by m , $\log(\det \Lambda)$, $\log(\tau)$ and $\log(1/\lambda_1^*)$. As $\log(\det \Lambda) \leq O(m \log(1/\lambda_1^*))$ we can just omit $\log(\det \Lambda)$. Making all quantities homogeneous with respect to lattice scaling, we obtain a classical complexity of $\text{poly}(m, \log \frac{a}{\lambda_1^*}, \log \frac{a}{\tau})$ bit operations.

The quantum complexity is driven by the Fourier transform in the dual lattice sampling and the Gaussian preparation step. Repeating the dual lattice sampling k times costs $O(kmQ \log(mQ))$ quantum gates and $O(mQ + n)$ qubits, where n is the number qubits required to store the values $|\mathbf{f}(x)\rangle$ of the quantum oracle in (see Theorem 3.7). Repeating k times the preparation of the Gaussian initial quantum state (within total variation distance $\eta = 1/k^2$) requires $O(kmQ \log(kmQ)^2)$ quantum gates and $O(mQ + \log(k)) = O(mQ)$ qubits (where we hide $\log(k)$ into $O(mQ)$), see Theorem 3.12. As discussed in Remark 3.13, the quantum gate complexity is slightly dominated by that of the Gaussian preparation step that occurs in Step 1 of Algorithm 2; it is $O(kmQ \log(kmQ)^2)$. The overall qubit complexity is $O(mQ + n)$.

For the estimation of the number of qubits Q needed ‘per dimension’, i.e., to prove Equation (3.24), we instantiate $\eta = 1/k^2$ and $\delta = 2^{-O(mk)} \cdot (\sqrt{m} \cdot a)^{-(m+1)} \cdot \det(\Lambda)^{-1} \cdot (\lambda_1^*)^2 \cdot \tau$ in Theorem 3.7 to obtain

$$\log(1/\delta) = (m+1) \log(\sqrt{m}a) + \log(\det(\Lambda)) + O(mk) - \log \tau - 2 \log(\lambda_1^*).$$

Noting that $m \log(\sqrt{m}a) + \log(\det(\Lambda)) \in O(k) \subseteq O(mk)$, we see that

$$O\left(\log \frac{a}{\eta \cdot \delta \lambda_1^*}\right) = O(mk) + O\left(\log \frac{1}{\lambda_1^* \cdot \tau}\right) + O(\log(a/\lambda_1^*))$$

Putting $O(m \log m)$ into $O(mk)$ in Equation (3.27) yields

$$Q = O(mk) + O\left(\log \frac{a}{\lambda_1^*}\right) + O\left(\log \frac{1}{\lambda_1^* \cdot \tau}\right), \quad (3.28)$$

□

3.4. Dual Lattice Sampling Algorithm

3.4.1. The Algorithm

Given a Λ -periodic function $\mathbf{f} : \mathbb{R}^m \rightarrow \mathcal{S}$ as discussed in Section 3.3, which maps a classical input x to a quantum state $|\mathbf{f}(x)\rangle$, we consider the following quantum algorithm (see Algorithm 2, or more graphically, Figure 3.4). The algorithm has oracle access to \mathbf{f} , meaning that it has access to a unitary that maps $|x\rangle|0\rangle$ to $|x\rangle|\mathbf{f}(x)\rangle$. As a matter of fact, we may assume the algorithm to have oracle access to a unitary that maps $|x\rangle|0\rangle$ to $|x\rangle|\mathbf{f}(Vx)\rangle$ for a parameter $V \in \mathbb{R}$ chosen by the algorithm. Per se, x may be arbitrary in \mathbb{R}^m ; for any concrete algorithm it is of course necessary to restrict x to some finite subset of \mathbb{R}^m .

The algorithm we consider follows the blueprint of the standard hidden-subgroup algorithm. Notable differences are that we need to discretize (and finitize) the continuous domain \mathbb{R}^m of the function, and the algorithm starts off with a superposition that is not uniform but follows a (discretized and finitized) Gaussian distribution. The reason for the latter choice is that Gaussian distributions decay very fast and behave nicely under the Fourier transform (as they are eigenfunctions of the Fourier transform).

The algorithm is given in Algorithm 2. It uses two quantum registers, each one consisting of a certain number of qubits. Associated to the first register are *grid points*: orthonormal bases $\{|x\rangle_{\mathbb{D}^m}\}_{x \in \mathbb{D}^m}$ and $\{|y\rangle_{\hat{\mathbb{D}}^m}\}_{y \in \hat{\mathbb{D}}^m}$ where the basis vectors are labeled by $x \in \mathbb{D}^m$ and $y \in \hat{\mathbb{D}}^m$, respectively, which we identify with elements $x \in \mathbb{D}_{\text{rep}}^m$ and $y \in \hat{\mathbb{D}}_{\text{rep}}^m$ (see Section 2.2.1). The second

3. The Continuous Hidden Subgroup Problem

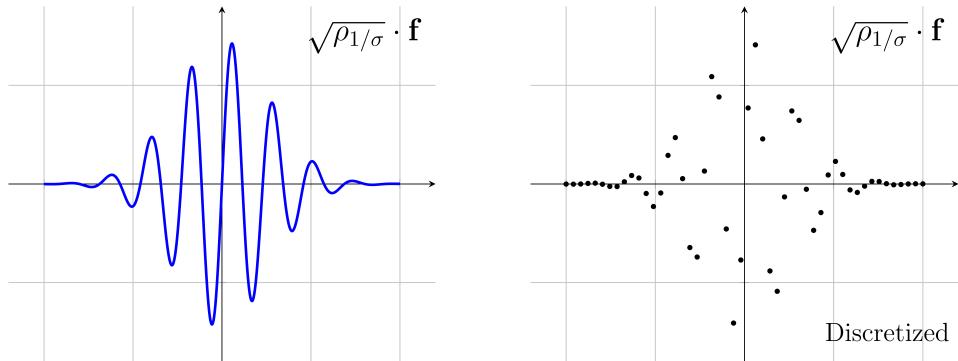


Figure 3.3.: Intuitively, it is easier to see the (quasi-)periodicity of the continuous signal (left) than that of the discrete signal (right). It is exactly the loss of information ‘between the sampling points’ that causes this chapter’s quantum algorithm to behave slightly erroneously or noisily. Of course, increasing the number of sampling points should reduce this noise; but it also causes the algorithm to need more expensive qubits. The analysis sought to keep the required qubits as low as possible, while still maintaining an acceptable error probability.

register has state space \mathcal{H} . The algorithm is parameterized by $q \in \mathbb{N}$ (which determines \mathbb{D}^m), $s > 0$ and $V > 0$.

Intuitively, the fraction $\frac{s}{V}$ is tightly related to the absolute precision of the output, whereas $\log q$ is connected with the number of qubits needed. In Algorithm 2, all quantum states described are *unnormalized* (i.e., do not have norm 1) but have all the same norm, due to the unitary operations in each step. In the analysis later, we show that, for adequately chosen parameters, the initial state $|\psi_0\rangle$, and therefore all states, are actually very *close* to normalized.

The description and analysis of Step 1 of Algorithm 2 is deferred to Appendix A.5. It will be shown (as summarized in Theorem 3.12) that its cost is comparable to the main cost of Algorithm 2, while contributing an error of at most $o(\eta)$ in the trace distance.

Algorithm 2: Quantum algorithm for the dual lattice sampling problem

1: **Prepare the Gaussian state**

$$|\psi_0\rangle := s^{m/2} \cdot \sum_{x \in \mathbb{D}^m} \sqrt{\rho_{1/s}(x)} \cdot |x\rangle_{\mathbb{D}^m} |0\rangle ;$$

2: **Apply the f -oracle**, yielding $s^{m/2} \cdot \sum_{x \in \mathbb{D}^m} \sqrt{\rho_{1/s}(x)} \cdot |x\rangle_{\mathbb{D}^m} |\mathbf{f}(Vx)\rangle ;$

3: **Apply the quantum Fourier transform on the first register**,
yielding the unnormalized state

$$s^{m/2} \cdot \sum_{x \in \mathbb{D}^m} \sum_{y \in \hat{\mathbb{D}}^m} \sqrt{\rho_{1/s}(x)} \cdot e^{-2\pi i \langle x, y \rangle} \cdot |y\rangle_{\hat{\mathbb{D}}^m} |\mathbf{f}(Vx)\rangle ;$$

4: **Measure the first register in the $\hat{\mathbb{D}}_{\text{rep}}^m$ -basis** yielding some $y \in \hat{\mathbb{D}}_{\text{rep}}^m$, and output $\frac{y}{V} ;$

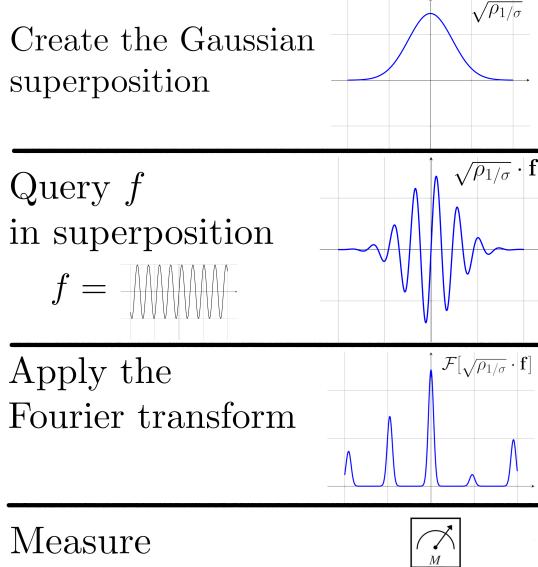


Figure 3.4.: A visual representation of Algorithm 2, if it would have been run on a ‘continuous’ quantum computer with infinitely many qubits. In reality, quantum computers have only finitely many qubits, leading to discretization errors. These errors are the main topic of this chapter. Note that the state after the Fourier transform ‘peaks’ at the dual lattice points.

3.4.2. The Figure of Merit

Recall that $N = \dim \mathcal{H} = 2^n$. Then the state after step (2) of Algorithm 2 equals, up to normalization,

$$|\psi\rangle := s^{m/2} \sum_{x \in \mathbb{D}^m} \sqrt{\rho_{1/s}(x)} |x\rangle_{\mathbb{D}^m} |\mathbf{f}(Vx)\rangle$$

which we can rewrite as

$$|\psi\rangle = \sum_{x \in \mathbb{D}^m} |x\rangle_{\mathbb{D}^m} |\mathbf{h}(x)\rangle$$

where

$$|\mathbf{h}(x)\rangle := s^{m/2} \sqrt{\rho_{1/s}(x)} \cdot |\mathbf{f}(Vx)\rangle.$$

Applying the quantum Fourier transform in step (3) maps this to

$$\begin{aligned} |\hat{\psi}\rangle &= q^{-m/2} \sum_{x \in \mathbb{D}^m} \sum_{y \in \hat{\mathbb{D}}^m} e^{-2\pi i \langle x, y \rangle} |y\rangle_{\hat{\mathbb{D}}^m} |\mathbf{h}(x)\rangle \\ &= q^{m/2} \sum_{y \in \hat{\mathbb{D}}^m} |y\rangle_{\hat{\mathbb{D}}^m} |\mathcal{F}_{\mathbb{D}^m} \{\mathbf{h}\}(y)\rangle, \end{aligned}$$

where the factor $q^{m/2}$ comes from the fact that, by our convention, the Fourier transform $\mathcal{F}_{\mathbb{D}^m}$ is scaled with the factor q^{-m} , while the quantum Fourier transform comes with a scaling factor $q^{-m/2}$.

Up to normalization, the probability to observe outcome $y \in \hat{\mathbb{D}}^m$ in step (4) thus is

$$\langle \hat{\psi} | (|y\rangle\langle y| \otimes \mathbb{I}) | \hat{\psi} \rangle = q^m \cdot \|\mathcal{F}_{\mathbb{D}^m} \{\mathbf{h}\}(y)\|_{\mathcal{H}}^2,$$

and so, for any ‘target’ subset $C \subset \hat{\mathbb{D}}^m$, the probability for the algorithm to produce an outcome $y \in C$ equals

$$\mathcal{D}(C) = \sum_{y \in C} \frac{\langle \hat{\psi} | (|y\rangle\langle y| \otimes \mathbb{I}) | \hat{\psi} \rangle}{\langle \psi_0 | \psi_0 \rangle} = \frac{\|1_C \cdot \mathcal{F}_{\mathbb{D}^m} \{\mathbf{h}\}\|_{\hat{\mathbb{D}}^m}^2}{\frac{s^m}{q^m} \sum_{x \in \mathbb{D}^m} \rho_{1/s}(x)}. \quad (3.29)$$

This target set are the points that one *would like to have as an outcome* after measuring. In our situation, this target set C consists of points close to dual lattice points ℓ^* , as those are considered ‘good’ measurement (see Figure 3.5).

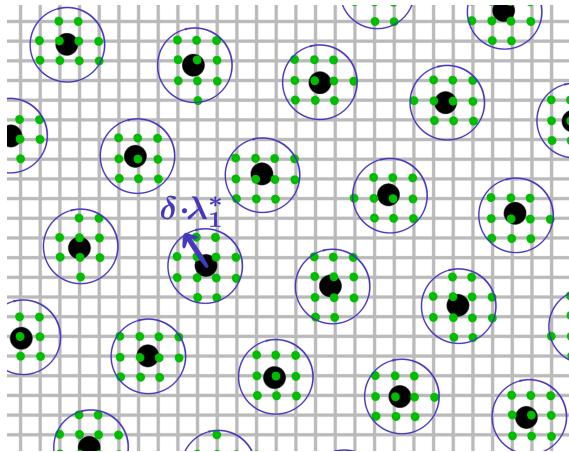


Figure 3.5.: The target set C consists of those grid points that are $\delta \cdot \lambda_1^*$ -close to the dual lattice Λ^* ; these points give valuable information about the dual lattice Λ^* . In this specific example, the target set consists of the green points and the blue circles around the black dual lattice points have radius $\delta \cdot \lambda_1^*$.

The algorithm's behavior in the limit

Intuitively, in the limit $q \rightarrow \infty$, the grid $\frac{1}{q}\mathbb{Z}^m$ becomes \mathbb{R}^m ; thus, neglecting constant factors, the function $\mathcal{F}_{\mathbb{D}^m}\{\mathbf{h}\}$ is expected to converge to

$$\mathcal{F}_{\mathbb{R}^m}\{\rho_{\sqrt{2}/s} \cdot \mathbf{f}(V \cdot)\} = \rho_{s/\sqrt{2}} \star \mathcal{F}_{\mathbb{R}^m}\{\mathbf{f}(V \cdot)\}.$$

Furthermore, when V is large enough compared to s , then, relative to the dual lattice $V\Lambda^*$, the Gaussian function behaves as a Dirac delta function. Thus, the above function is then supported by $V\Lambda^*$ and takes on the values $|c_{\ell^*}\rangle$. Hence, by taking square norms, we get the claimed $\langle c_{\ell^*} | c_{\ell^*} \rangle$.

Below, we prove that this intuition is indeed correct, and we work out the actual “rate of convergence”.

3.5. Analysis

3.5.1. Proof Overview

In the following few paragraphs we give an overview of the proof of correctness of Algorithm 2. The main idea boils down to showing that the finite Fourier transform is close to the continuous Fourier transform on the function $\mathbf{h} = \mathbf{f} \cdot \rho_{1/s}$. They are indeed close due to the smoothness of the Gaussian and the Lipschitz-continuity of the oracle function \mathbf{f} .

The unnormalized initial state $|\psi_0\rangle$ has approximately norm one. By the smoothing argument of Banaszczyk, we derive that the initial state's norm satisfies $\langle \psi_0 | \psi_0 \rangle = \frac{s^m}{q^m} \sum_{x \in \mathbb{D}^m} \rho_{1/s}(x) \approx 1$. So, the initial state might not be *perfectly* normalized, but it is almost. Therefore,

$$\mathcal{D}(C) = \frac{\|1_C \cdot \mathcal{F}_{\mathbb{D}^m} \{\mathbf{h}\}\|_{\mathbb{D}^m}^2}{\frac{s^m}{q^m} \sum_{x \in \mathbb{D}^m} \rho_{1/s}(x)} \approx \|1_C \cdot \mathcal{F}_{\mathbb{D}^m} \{\mathbf{h}\}\|_{\mathbb{D}^m}^2$$

meaning that we can focus on the latter quantity, that consists just of the norm of the Fourier transformed function \mathbf{h} .

Replacing the function \mathbf{h} by its \mathbb{T}^m -periodization $\mathbf{h}|^{\mathbb{T}^m}$. The function $\mathbf{h} = s^{m/2} \cdot \mathbf{f} \cdot \rho_{\sqrt{2}/s}$ is a product of the function \mathbf{f} and a Gaussian that is narrow enough to be contained within the centered unit cube. Therefore, periodization of \mathbf{h} with respect to the unit cube $[-\frac{1}{2}, \frac{1}{2}]^m$ (i.e., the central representative of the unit torus) doesn't differ too much from restricting \mathbf{h} to the torus. Therefore,

$$\|1_C \cdot \mathcal{F}_{\mathbb{D}^m} \{\mathbf{h}\}\|_{\mathbb{D}^m}^2 \approx \|1_C \cdot \mathcal{F}_{\mathbb{D}^m} \{\mathbf{h}|^{\mathbb{T}^m}\}\|_{\mathbb{D}^m}^2.$$

Replacing the finite \mathbb{D}^m -Fourier transform by the \mathbb{T}^m -Fourier transform. Because the function \mathbf{h} is Lipschitz-continuous, changing the finite Fourier transform into a continuous one over the torus \mathbb{T}^m gives us a error that

depends mainly on the discretization parameter q and the Lipschitz constant $\text{Lip}(\mathbf{f})$.

$$\|1_C \cdot \mathcal{F}_{\mathbb{D}^m}\{\mathbf{h}|^{\mathbb{T}^m}\}\|_{\mathbb{D}^m}^2 \approx \|1_C \cdot \mathcal{F}_{\mathbb{T}^m}\{\mathbf{h}|^{\mathbb{T}^m}\}\|_{\mathbb{Z}^m}^2.$$

Replacing the \mathbb{T}^m -Fourier transform by the \mathbb{R}^m -Fourier transform. Using the Poisson summation formula, one can derive an equality between the Fourier transform of $\mathbf{h}|^{\mathbb{T}^m}$ over the torus \mathbb{T}^m and the Fourier transform of \mathbf{h} over the reals \mathbb{R}^m .

$$\|1_C \cdot \mathcal{F}_{\mathbb{T}^m}\{\mathbf{h}|^{\mathbb{T}^m}\}\|_{\mathbb{Z}^m}^2 = \|1_C \cdot \mathcal{F}_{\mathbb{R}^m}\{\mathbf{h}\}\|_{\mathbb{Z}^m}^2.$$

Relating the \mathbb{R}^m -Fourier transform with the Fourier coefficients $|c_{\ell^}\rangle$ of $|\mathbf{f}\rangle$.* As \mathbf{h} is essentially a product of \mathbf{f} and a relatively wide Gaussian, one can apply the convolution theorem to obtain the real Fourier transform of \mathbf{h} . This Fourier transform is then very much related with the Fourier coefficients $|c_{\ell^*}\rangle$ of \mathbf{f} .

$$\|1_C \cdot \mathcal{F}_{\mathbb{R}^m}\{\mathbf{h}\}\|_{\mathbb{Z}^m}^2 \approx \sum_{\ell^* \in \Lambda^*} \langle c_{\ell^*} | c_{\ell^*} \rangle \cdot \iota_C(\ell^*)$$

The function ι_C here acts as sort-of an indicator function; one can think of $\iota_C(\ell^*)$ being close to one whenever ℓ^* is in the ‘target set’ C and zero otherwise. Recall that this target set are the ‘wanted’ points, i.e., the desired outcomes after measuring the quantum state. In our situation, this target set C consists of points close $\delta\lambda_1^*$ -close to dual lattice points ℓ^* , as those are considered ‘good’ measurements; they namely give valuable information about the dual lattice Λ^* .

Lower bounding the success probability by means of Fourier coefficients of \mathbf{f} . In particular, one can show that, up to a small error, the function ι_C indeed acts as an indicator function. Whenever a large enough ball around a dual lattice point ℓ^* is contained in C , the value of $\iota_C(\ell^*)$ approximates one.

$$\mathcal{D}(C) \approx \sum_{\ell^* \in \Lambda^*} \langle c_{\ell^*} | c_{\ell^*} \rangle \cdot \iota_C(\ell^*) \geq \sum_{\substack{\ell^* \in \Lambda^* \\ \mathcal{B}_{\delta\lambda_1^*}(\ell^*) \cap \mathbb{Z}^m \subseteq C}} \langle c_{\ell^*} | c_{\ell^*} \rangle. \quad (3.30)$$

3. The Continuous Hidden Subgroup Problem

Taking into account the bounded output of Algorithm 2 and finalizing the analysis. The output distribution \mathcal{D} of Algorithm 2 has support only in $[-q/2, q/2]^m$. So, for any $S \subseteq \Lambda^*$ the probability p_S from Problem 3.6 applied to the output distribution of Algorithm 2 satisfies

$$\begin{aligned} p_S &= \mathcal{D}(\mathcal{B}_{\delta\lambda_1^*}(S)) = \mathcal{D}(\mathcal{B}_{\delta\lambda_1^*}(S) \cap [-q/2, q/2]^m) \\ &\gtrapprox \sum_{\substack{\ell^* \in \Lambda^* \\ \ell^* \in S \cap [-q/2, q/2]^m}} \langle c_{\ell^*} | c_{\ell^*} \rangle \gtrapprox \sum_{\ell^* \in S} \langle c_{\ell^*} | c_{\ell^*} \rangle \end{aligned}$$

where the first ‘approximate inequality’ (which is an inequality up to some small error) is obtained from Equation (3.30) and the last ‘approximate inequality’ holds by the fact that the ‘tail’ of the Fourier coefficients of \mathbf{f} has small weight, i.e., $\sum_{|\ell^*| > q/2} \langle c_{\ell^*} | c_{\ell^*} \rangle$ is small.

Summarizing, this error mainly occurs because of the phrasing of the Problem 3.6. It makes the suggestion that the distribution \mathcal{D} should have unbounded support and should be able to reach any dual lattice point, whereas in reality (for the output distribution of Algorithm 2) this is very much not the case. The error induced by this discrepancy is, as a consequence, essentially the combined weight (i.e., the ‘lost probability’) of the lattice points unreachable by the output distribution of Algorithm 2.

The velocity parameter V . In the formal analysis below, we sometimes temporarily assume that the velocity parameter equals one, i.e., $V = 1$. This is for sake of clarity and can be done without loss of generality, since for arbitrary V the very same reasoning can be applied to the function $\mathbf{f}_V := \mathbf{f}(V \cdot)$. This affects the quantities involved in the sense that Λ^* becomes $V\Lambda^*$, λ_1^* becomes $V \cdot \lambda_1^*$ and $\text{Lip}(\mathbf{f}_V)$ becomes $V \text{Lip}(\mathbf{f})$.

To be clear, the end results and errors involved are always stated *for general V* . Moreover, whenever the assumption $V = 1$ occurs in a proof or a line of reasoning, we will always explicitly say so, in order to avoid confusion.

3.5.2. Formal Analysis

The unnormalized initial state $|\psi_0\rangle$ has approximately norm one

By the smoothing lemma (see Lemma 2.31), we have, whenever $q/s \geq \sqrt{m}$,

$$\begin{aligned}\langle \psi_0 | \psi_0 \rangle &= \frac{s^m}{q^m} \sum_{x \in \mathbb{D}^m} \rho_{1/s}(x) \leq \frac{s^m}{q^m} \cdot \rho_{1/s} \left(\frac{1}{q} \mathbb{Z}^m \right) \leq 1 + 2\beta_{q/s} \\ &\leq 1 + O(e^{-q^2/s^2}).\end{aligned}$$

Therefore,

$$\left| \frac{\|1_C \cdot \mathcal{F}_{\mathbb{D}^m} \{ \mathbf{h} \}\|_{\hat{\mathbb{D}}^m}^2}{\frac{s^m}{q^m} \sum_{x \in \mathbb{D}^m} \rho_{1/s}(x)} - \|1_C \cdot \mathcal{F}_{\mathbb{D}^m} \{ \mathbf{h} \}\|_{\hat{\mathbb{D}}^m}^2 \right| \leq O(e^{-q^2/s^2}). \quad (3.31)$$

By requiring that $q/s \geq \sqrt{m + \log(\eta^{-1})}$, we can safely neglect this error.

Replacing the function \mathbf{h} by its \mathbb{T}^m -periodization $\mathbf{h}|^{\mathbb{T}^m}$

By the linearity of the Fourier transform, by the fact that 1_C is an indicator function and by Parseval's theorem, one can deduce

$$\begin{aligned}\|1_C \cdot \mathcal{F}_{\mathbb{D}^m} \{ \mathbf{h} \} - 1_C \cdot \mathcal{F}_{\mathbb{D}^m} \{ \mathbf{h}|^{\mathbb{T}^m} \}\|_{\hat{\mathbb{D}}^m} &\leq \|\mathcal{F}_{\mathbb{D}^m} \{ \mathbf{h} - \mathbf{h}|^{\mathbb{T}^m} \}\|_{\hat{\mathbb{D}}^m} \\ &= \|\mathbf{h}|^{\mathbb{T}^m} - \mathbf{h}\|_{\mathbb{D}^m}.\end{aligned}$$

Writing out the definition of the functions $\mathbf{h} = s^{m/2} \cdot \mathbf{f} \cdot \rho_{\sqrt{2}/s}$ and $\mathbf{h}|^{\mathbb{T}^m} = \sum_{z \in \mathbb{Z}^m} \mathbf{h}(z + \cdot)$, we obtain

$$\begin{aligned}\|\mathbf{h}|^{\mathbb{T}^m} - \mathbf{h}\|_{\mathbb{D}^m}^2 &= \frac{1}{q^m} \sum_{x \in \mathbb{D}^m} \left\| \sum_{z \in \mathbb{Z}^m \setminus 0} \mathbf{h}(x + z) \right\|_{\mathcal{H}}^2 \\ &\leq \frac{s^m}{q^m} \sum_{x \in \mathbb{D}^m} \left(\sum_{z \in \mathbb{Z}^m \setminus 0} \rho_{\sqrt{2}/s}(x + z) \cdot \|\mathbf{f}(V(x + z))\|_{\mathcal{H}} \right)^2.\end{aligned}$$

3. The Continuous Hidden Subgroup Problem

Since $\|\mathbf{f}(x)\|_{\mathcal{H}} = \sqrt{\langle \mathbf{f}(x) | \mathbf{f}(x) \rangle} = 1$, as $|\mathbf{f}(x)\rangle$ is a quantum state for any $x \in \mathbb{R}^m$, above expression is bounded by

$$\begin{aligned} \frac{s^m}{q^m} \sum_{x \in \mathbb{D}^m} \left(\underbrace{\sum_{z \in \mathbb{Z}^m \setminus \{0\}} \rho_{\sqrt{2}/s}(x+z)}_{\leq 2 \cdot \beta_{\frac{s}{2\sqrt{2}}}} \right)^2 &\leq \frac{s^m \cdot |\mathbb{D}^m|}{q^m} \cdot (2 \cdot \beta_{\frac{s}{2\sqrt{2}}})^2 \\ &\leq 4 \cdot s^m \cdot (\beta_{\frac{s}{2\sqrt{2}}})^2, \end{aligned}$$

as $\rho_{\sqrt{2}/s}(\mathbb{Z}^m \setminus \{0\} + x) \leq 2 \cdot \beta_{\frac{s}{2\sqrt{2}}}$, from Banaszczyk's tail bound in Corollary 2.30. By the reverse triangle inequality, provided that $s \geq \sqrt{8m}$, we conclude

$$\left| \|1_C \cdot \mathcal{F}_{\mathbb{D}^m} \{ \mathbf{h} \}\|_{\hat{\mathbb{D}}^m}^2 - \|1_C \cdot \mathcal{F}_{\mathbb{D}^m} \{ \mathbf{h} | \mathbb{T}^m \}\|_{\hat{\mathbb{D}}^m}^2 \right| \leq O(s^m e^{-s^2/8}). \quad (3.32)$$

By requiring that $s \geq \sqrt{8m \log(m) + \log(\eta^{-1})}$, we can safely neglect this error.

Replacing the finite \mathbb{D}^m -Fourier transform by the \mathbb{T}^m -Fourier transform

Using Theorem 2.8 with $\mathbf{h} | \mathbb{T}^m$, one obtains

$$\left| \|1_C \cdot \mathcal{F}_{\mathbb{D}^m} \{ \mathbf{h} | \mathbb{T}^m \}\|_{\hat{\mathbb{D}}^m} - \|1_C \cdot \mathcal{F}_{\mathbb{T}^m} \{ \mathbf{h} | \mathbb{T}^m \}\|_{\mathbb{Z}^m} \right| \leq \frac{4\pi\sqrt{m} \text{Lip}(\mathbf{h} | \mathbb{T}^m)}{q} \quad (3.33)$$

$$\leq O \left(\frac{\sqrt{m} s^{m/2} (V \text{Lip}(\mathbf{f}) + s^2)}{q} \right). \quad (3.34)$$

Remark 3.14. In above inequality the indicator function 1_C is used as a function on both \mathbb{D}^m and \mathbb{Z}^m . The function 1_C on \mathbb{Z}^m must be interpreted as having the same values on $\mathbb{D}_{\text{rep}}^m \subseteq \mathbb{Z}^m$ as on \mathbb{D}^m and having value zero otherwise.

Lemma 3.15. *Assume that $s \geq 4\sqrt{m}$. Then, for the Lipschitz constant $\text{Lip}(\mathbf{h}|^{\mathbb{T}^m})$ of $\mathbf{h}|^{\mathbb{T}^m}$ holds*

$$\text{Lip}(\mathbf{h}|^{\mathbb{T}^m}) \leq s^{m/2} (2V \text{Lip}(\mathbf{f}) + \pi s^2).$$

Proof. For the sake of clarity, we assume $V = 1$ throughout this proof; at the end we will then have to replace $\text{Lip}(\mathbf{f})$ by $V \text{Lip}(\mathbf{f})$. Also, we will temporarily omit the constant term $s^{m/2}$ in the definition of \mathbf{h} and use ρ for $\rho_{\sqrt{2}/s}$; thus calculating with $\mathbf{h} = \mathbf{f} \cdot \rho$ instead. In the final step, the multiplicative term $s^{m/2}$ will then be multiplied again to the end result.

By applying the triangle inequality multiple times, using the fact that $\|\mathbf{f}(x)\|_{\mathcal{H}} = 1$ for all $x \in \mathbb{R}^m$ and using the Lipschitz-continuity of \mathbf{f} , one obtains, for every $x, y \in \mathbb{R}^m$,

$$\begin{aligned} \|\mathbf{h}(x) - \mathbf{h}(y)\|_{\mathcal{H}} &\leq \|\mathbf{f}(x)(\rho(x) - \rho(y))\|_{\mathcal{H}} + \|(\mathbf{f}(x) - \mathbf{f}(y))\rho(y)\|_{\mathcal{H}} \\ &\leq |\rho(x) - \rho(y)| + \text{Lip}(\mathbf{f}) \cdot \|x - y\|_{\mathbb{R}^m} \cdot \rho(y) \end{aligned} \quad (3.35)$$

By periodizing with respect to the unit torus $\mathbb{T}^m = \mathbb{R}^m / \mathbb{Z}^m$ and applying the triangle inequality, we obtain, for all $x, y \in [-1/2, 1/2]^m$,

$$\begin{aligned} \|\mathbf{h}|^{\mathbb{T}^m}(x) - \mathbf{h}|^{\mathbb{T}^m}(y)\|_{\mathcal{H}} &\leq \sum_{z \in \mathbb{Z}^m} |\rho(x + z) - \rho(y + z)| \\ &\quad + \text{Lip}(\mathbf{f}) \cdot \|x - y\|_{\mathbb{T}^m} \cdot \sum_{z \in \mathbb{Z}^m} \rho(y + z) \end{aligned} \quad (3.36)$$

By smoothing arguments of Banaszczyk, one deduces that $\rho_{\sqrt{2}/s}(y + \mathbb{Z}^m) \leq 2$ (see Corollary 2.30), where we use the assumption $s \geq 4\sqrt{m}$. By the reasoning in Lemma A.33, we have that

$$\begin{aligned} &\sum_{z \in \mathbb{Z}^m} |\rho_{\sqrt{2}/s}(x + z) - \rho_{\sqrt{2}/s}(y + z)| \\ &\leq \pi s^2 / 2 \cdot \|x - y\|_{\mathbb{T}^m} \underbrace{\sum_{z \in \mathbb{Z}^m} \rho_{\sqrt{2}/s}(x + y + 2z) \|x + y + 2z\|}_{\leq 2} \\ &\leq \pi s^2 \cdot \|x - y\|_{\mathbb{T}^m}, \end{aligned} \quad (3.37)$$

3. The Continuous Hidden Subgroup Problem

where the last inequality can be obtained by absorbing $\|x + y + 2z\|$ into the Gaussian and applying smoothing arguments again; $\rho_{1/s}(x) \cdot \|x\| \leq \rho_{2/s}(x)$ for all $x \in \mathbb{R}^m$ and $s \geq \sqrt{m}$, and $\rho_{\sqrt{8}/s}(\mathbb{Z}^m) \leq \rho_{\sqrt{m}}(\mathbb{Z}^m) \leq 1 + 2 \cdot \beta_{\sqrt{m}} \leq 2$, for $s \geq 4\sqrt{m}$ (see Lemma 2.29). In other words,

$$\begin{aligned} \sum_{z \in \mathbb{Z}^m} \rho_{\sqrt{8}/s}(x + y + 2z) \|x + y + 2z\| &\leq \sum_{z \in \mathbb{Z}^m} \rho_{\sqrt{32}/s}(x + y + 2z) \\ &\leq \rho_{\sqrt{32}/s}(2 \cdot \mathbb{Z}^m) = \rho_{\sqrt{8}/s}(\mathbb{Z}^m) \leq 2. \end{aligned}$$

By combining Equations (3.35) to (3.37), multiplying the factor $s^{m/2}$ and replacing $\text{Lip}(\mathbf{f})$ by $V \cdot \text{Lip}(\mathbf{f})$ we obtain the final result. \square

Replacing the \mathbb{T}^m -Fourier transform by the \mathbb{R}^m -Fourier transform

Apply the Poisson summation formula (see Corollary 2.5) to conclude that

$$\|1_C \cdot \mathcal{F}_{\mathbb{T}^m}\{\mathbf{h}\}^{\mathbb{T}^m}\|_{\mathbb{Z}^m} = \|1_C \cdot \mathcal{F}_{\mathbb{R}^m}\{\mathbf{h}\}\|_{\mathbb{Z}^m},$$

where $\mathcal{F}_{\mathbb{R}^m}\{\mathbf{h}\}$ is temporarily identified with its restriction to \mathbb{Z}^m .

Relating the \mathbb{R}^m -Fourier transform with the Fourier coefficients $|c_{\ell^*}\rangle$ of $|\mathbf{f}\rangle$

By applying the convolution theorem as outlined in Equation (2.9) of Section 2.2.2, we see that

$$\begin{aligned} \mathcal{F}_{\mathbb{R}^m}\{\mathbf{h}\}[y] &= \mathcal{F}_{\mathbb{R}^m/\Lambda}\{\mathbf{f}(V \cdot)\} \star \mathcal{F}_{\mathbb{R}^m}\{s^{m/2} \rho_{\sqrt{2}/s}(\cdot)\}(y) \\ &= \left(\frac{2}{s}\right)^{m/2} \sum_{\ell^* \in \Lambda^*} |c_{\ell^*}\rangle \rho_{s/\sqrt{2}}(y - V\ell^*), \end{aligned}$$

where $|c_{\ell^*}\rangle$ are the vectorial Fourier coefficients of \mathbf{f} . Therefore,

$$\begin{aligned} &\|\mathcal{F}_{\mathbb{R}^m}\{\mathbf{h}\}[y]\|_{\mathcal{H}}^2 \\ &= \left(\frac{2}{s}\right)^m \sum_{k^* \in \Lambda^*} \sum_{\ell^* \in \Lambda^*} \langle c_{\ell^*} | c_{k^*} \rangle \rho_{s/\sqrt{2}}(y - V\ell^*) \rho_{s/\sqrt{2}}(y - Vk^*) \\ &= \left(\frac{2}{s}\right)^m \sum_{u^* \in \frac{1}{2}\Lambda^*} \sum_{v^* \in u^* + \Lambda^*} \langle c_{v^*+u^*} | c_{v^*-u^*} \rangle \rho_{s/2}(Vu^*) \rho_{s/2}(y - Vv^*), \end{aligned} \quad (3.38)$$

where the latter is obtained by the variable substitution $u^* = \frac{\ell^* - k^*}{2}$, $v^* = \frac{\ell^* + k^*}{2}$, and using the multiplicative properties of Gaussian functions (see Lemma 2.23), like $\rho_{s/\sqrt{2}}(x)\rho_{s/\sqrt{2}}(y) = \rho_{s/2}((x+y)/2)\rho_{s/2}((x-y)/2)$ for all $x, y \in \mathbb{R}^m$.

Definition 3.16. For any subset $C \subseteq \mathbb{Z}^m$, any $s > 0$ and any $\ell^* \in \Lambda^*$, we define $\iota_C : \Lambda^* \rightarrow \mathbb{R}_{>0}$ by the following rule,

$$\iota_C(\ell^*) := \left(\frac{2}{s}\right)^m \sum_{y \in C} \rho_{s/2}(y - V\ell^*),$$

where leave out the dependence on s and V in the notation.

The above definition of ι_C is mainly to make the notation in this analysis more compact. But this function on Λ^* also has an intuitive interpretation; it is the cumulative Gaussian weight of all points in C around ℓ^* (or, $V \cdot \ell^*$ in the case of scaling with V). So, if C contains many close points around ℓ^* (see Figure 3.5 and Figure 3.6), this cumulative Gaussian weight approaches 1, whereas if there are no points in C around ℓ^* , this weight approaches zero. Summarizing, the value $\iota_C(\ell^*)$ quantifies the number of close points around ℓ^* ; a value of 1 indicates many good close points in C , whereas a value near 0 indicates no good close points (see Figure 3.6).

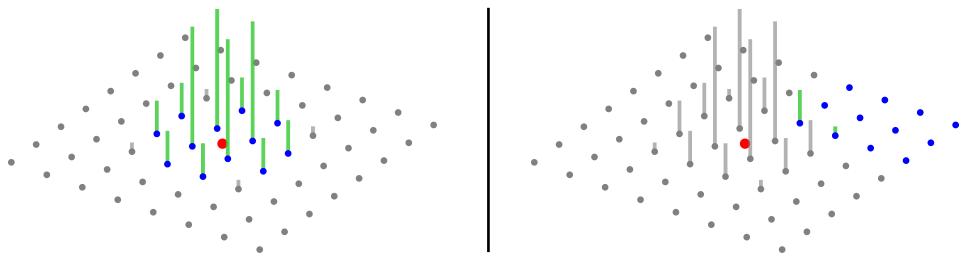


Figure 3.6.: The function $\iota_C(\ell^*)$ equals the cumulative Gaussian weight of all points in C around ℓ^* . In the left panel above, the set C contains many points around the red lattice point ℓ^* , yielding a cumulative Gaussian weight approaching one, i.e., $\iota_C(\ell^*) \approx 1$. In the right panel, set C only contains a few points close to the lattice point, yielding a very low Gaussian weight, i.e., $\iota_C(\ell^*) \approx 0$.

3. The Continuous Hidden Subgroup Problem

Lemma 3.17. *Let $V, s > 0$ satisfy the conditions $V\lambda_1^*/s \geq \sqrt{m}$ and $s \geq \sqrt{m}$. Then, for any $C \subseteq [q]_c^m$, we have*

$$\left| \|1_C \mathcal{F}_{\mathbb{R}^m}\{\mathbf{h}\}\|_{\mathbb{Z}^m}^2 - \sum_{\ell^* \in \Lambda^*} \langle c_{\ell^*} | c_{\ell^*} \rangle \iota_C(\ell^*) \right| \leq O(e^{-(V\lambda_1^*/s)^2}). \quad (3.39)$$

Proof. Without loss of generality, we assume in the rest of the proof that $V = 1$, as sketched in the last paragraph of Section 3.5.1. At the end of the proof we will then replace λ_1^* by $V \cdot \lambda_1^*$.

By writing out the definition of the norm over \mathbb{Z}^m and using Equation (3.38), we obtain

$$\begin{aligned} \|1_C \mathcal{F}_{\mathbb{R}^m}\{\mathbf{h}\}\|_{\mathbb{Z}^m}^2 &= \sum_{y \in C} \|\mathcal{F}_{\mathbb{R}^m}\{\mathbf{h}\}[y]\|_{\mathcal{H}}^2 \\ &= \left(\frac{2}{s}\right)^m \sum_{y \in C} \sum_{\substack{u^* \in \frac{1}{2}\Lambda^* \\ v^* \in u^* + \Lambda^*}} \langle c_{v^*+u^*} | c_{v^*-u^*} \rangle \rho_{s/2}(u^*) \rho_{s/2}(y - v^*). \end{aligned}$$

By swapping the summation over C to the right, we deduce

$$\|1_C \mathcal{F}_{\mathbb{R}^m}\{\mathbf{h}\}\|_{\mathbb{Z}^m}^2 = \sum_{\substack{u^* \in \frac{1}{2}\Lambda^* \\ v^* \in u^* + \Lambda^*}} \langle c_{v^*+u^*} | c_{v^*-u^*} \rangle \rho_{s/2}(u^*) \iota_C(v^*).$$

We split above sum into a part where $u^* = 0$ and a part where $u^* \neq 0$. Notice that for the case $u^* = 0$, the inner product $\langle c_{v^*+u^*} | c_{v^*-u^*} \rangle$ becomes $\langle c_{v^*} | c_{v^*} \rangle$ and $\rho_{s/2}(u^*) = 1$. This yields

$$\begin{aligned} \|1_C \mathcal{F}_{\mathbb{R}^m}\{\mathbf{h}\}\|_{\mathbb{Z}^m}^2 &= \sum_{\ell^* \in \Lambda^*} \langle c_{\ell^*} | c_{\ell^*} \rangle \cdot \iota_C(\ell^*) \\ &\quad + \sum_{u^* \in \frac{1}{2}\Lambda^* \setminus \{0\}} \rho_{s/2}(u^*) \sum_{v^* \in u^* + \Lambda^*} \langle c_{v^*+u^*} | c_{v^*-u^*} \rangle \cdot \iota_C(v^*). \end{aligned} \quad (3.40)$$

In order to achieve the claim of this lemma, it is enough to bound the second term (where $u^* \neq 0$) in Equation (3.40). As we assumed that $s \geq \sqrt{m}$, we can bound $\iota_C(v^*) \leq \left(\frac{2}{s}\right)^m \rho_{s/2}(\mathbb{Z}^m + t) \leq 2$ for any $v^* \in \mathbb{R}^m$ and $C \subseteq \mathbb{Z}^m$ by applying smoothing arguments (see Corollary 2.32). The sum of the ‘shifted

inner products' of the Fourier coefficients is bounded by one, as can be seen by applying the Cauchy-Schwarz inequality and the inequality of arithmetic and geometric means.

$$\begin{aligned}
 \left| \sum_{v^* \in u^* + \Lambda^*} \langle c_{v^*+u^*} | c_{v^*-u^*} \rangle \right| &\leq \sum_{v^* \in \Lambda^*} \sqrt{\langle c_{v^*+2u^*} | c_{v^*+2u^*} \rangle \langle c_{v^*} | c_{v^*} \rangle} \\
 &\leq \sum_{v^* \in \Lambda^*} \frac{\langle c_{v^*+2u^*} | c_{v^*+2u^*} \rangle + \langle c_{v^*} | c_{v^*} \rangle}{2} = \|\mathbf{f}\|_{\mathbb{R}^m/\Lambda}^2 = 1.
 \end{aligned}$$

Combining above reasoning with a tail bound of Banaszczyk (Lemma 2.29) the $u^* \neq 0$ part in Equation (3.40) can be bounded as follows.

$$\begin{aligned}
 &\sum_{u^* \in \frac{1}{2}\Lambda^* \setminus 0} \rho_{s/2}(u^*) \left| \sum_{v^* \in u^* + \Lambda^*} \langle c_{v^*+u^*} | c_{v^*-u^*} \rangle \cdot \underbrace{\iota_C(v^*)}_{\leq 2} \right| \\
 &\leq 2 \sum_{u^* \in \frac{1}{2}\Lambda^* \setminus 0} \rho_{s/2}(u^*) \underbrace{\left| \sum_{v^* \in u^* + \Lambda^*} \langle c_{v^*+u^*} | c_{v^*-u^*} \rangle \right|}_{\leq 1} \\
 &\leq 2 \sum_{u^* \in \frac{1}{2}\Lambda^* \setminus 0} \rho_{s/2}(u^*) \leq 2 \cdot \rho_s(\Lambda^* \setminus 0) \leq 4 \cdot \beta_{\lambda_1^*/s}.
 \end{aligned}$$

In order to drop the assumption that $V = 1$ from the start of the proof, we need to replace λ_1^* by $V \cdot \lambda_1^*$ in above expression. Applying the bound $4 \cdot \beta_{V\lambda_1^*/s} \leq O(e^{-(V\lambda_1^*/s)^2})$ for $V\lambda_1^*/s \geq \sqrt{m}$ yields the final claim. \square

By requiring that $V\lambda_1^*/s \geq \sqrt{m + \log(\eta^{-1})}$, we can safely neglect the error from Lemma 3.17.

3. The Continuous Hidden Subgroup Problem

Lower bounding the success probability by means of Fourier coefficients of \mathbf{f}

Whenever $\mathcal{B}_{\delta\lambda_1^*V}(V\ell^*) \cap \mathbb{Z}^m \subseteq C$ for an $\ell^* \in \Lambda^*$, it holds that

$$\begin{aligned}\iota_C(\ell^*) &= \left(\frac{2}{s}\right)^m \sum_{y \in C} \rho_{s/2}(y - V\ell^*) \geq \left(\frac{2}{s}\right)^m \sum_{y \in \mathcal{B}_{V\delta\lambda_1^*}(V\ell^*) \cap \mathbb{Z}^m} \rho_{s/2}(y - V\ell^*) \\ &\geq \left(\frac{2}{s}\right)^m \rho_{s/2}(\mathbb{Z}^m) \left(1 - \beta_{2V\delta\lambda_1^*/s}\right) \geq (1 - 2 \cdot \beta_{s/2})(1 - \beta_{2V\delta\lambda_1^*/s}),\end{aligned}$$

where the second inequality follows from Banaszczyk's tail bound (see Lemma 2.25) and the last from the smoothing bound in Lemma 2.31. In other words, $\iota_C(\ell^*)$ is close to one if C contains all vectors in $\hat{\mathbb{D}}^m$ that are $\delta\lambda_1^*V$ -close to $V\ell^*$. This coincides with the intuitive explanation after Definition 3.16. Note that $\delta\lambda_1^*$ is the maximum distance from a dual lattice point ℓ^* required to consider the output valuable.

It follows then that

$$\begin{aligned}\left| \sum_{\ell^* \in \Lambda^*} \langle c_{\ell^*} | c_{\ell^*} \rangle \iota_C(\ell^*) - \sum_{\substack{\ell^* \in \Lambda^* \\ \mathcal{B}_{V\delta\lambda_1^*}(V\ell^*) \cap \mathbb{Z}^m \subseteq C}} \langle c_{\ell^*} | c_{\ell^*} \rangle \right| \\ \leq O(e^{-s^2/4}) + O(e^{-(2V\delta\lambda_1^*/s)^2}),\end{aligned}\tag{3.41}$$

where we use the fact that $\sum_{\ell^* \in \Lambda^*} \langle c_{\ell^*} | c_{\ell^*} \rangle = \|\mathbf{f}\|_{\mathbb{R}^m/\Lambda}^2 = 1$. By requiring that $\delta V\lambda_1^*/s \geq \sqrt{m + \log(\eta^{-1})}$ and $s \geq 4\sqrt{m + \log(\eta^{-1})}$, we can safely neglect this error.

Taking into account the bounded output of Algorithm 2 and finalizing the analysis

As the output distribution \mathcal{D} of Algorithm 2 has support only in $[-q/2, q/2]^m$, we have, for any $S \subseteq \Lambda^*$,

$$\mathcal{D}(\mathcal{B}_{\delta\lambda_1^*}(S)) = \mathcal{D}(\mathcal{B}_{\delta\lambda_1^*}(S) \cap [-q/2, q/2]^m).$$

By simply splitting the set $S \subseteq \Lambda^*$ into an ‘tail part’ $S_{\text{tail}} = S \setminus [-q/4, q/4]^m$ and a ‘bounded, finite part’ $S_{\text{fin}} = S \cap [-q/4, q/4]^m$, we obtain

$$\begin{aligned} \sum_{\ell^* \in S} \langle c_{\ell^*} | c_{\ell^*} \rangle - \mathcal{D}(\mathcal{B}_{\delta\lambda_1^*}(S)) &= \underbrace{\sum_{\ell^* \in S_{\text{tail}}} \langle c_{\ell^*} | c_{\ell^*} \rangle}_{\text{Small because of a tail bound}} \\ &+ \underbrace{\sum_{\ell^* \in S_{\text{fin}}} \langle c_{\ell^*} | c_{\ell^*} \rangle - \mathcal{D}(\mathcal{B}_{\delta\lambda_1^*}(S) \cap [-q/2, q/2]^m)}_{\text{Small because of the error analysis}}. \end{aligned} \quad (3.42)$$

By the fact that \mathbf{f} is a Lipschitz continuous function, its Fourier coefficients have a tail bound. By applying Corollary 2.34 with $B = q/4$, we obtain the following bound

$$\sum_{\ell^* \in S_{\text{tail}}} \langle c_{\ell^*} | c_{\ell^*} \rangle \leq \sum_{\ell^* \in \Lambda^* \setminus [-q/4, q/4]^m} \langle c_{\ell^*} | c_{\ell^*} \rangle \leq \frac{4 \cdot \text{Lip}(\mathbf{f})^2}{\pi^2 q^2}.$$

The summand in Equation (3.42) is, by the full error analysis, bounded by

$$\begin{aligned} \sum_{\mathcal{B}_{\delta\lambda_1^*}(\ell^*) \subseteq \mathcal{B}_{\delta\lambda_1^*}(S) \cap [-q/2, q/2]^m} \langle c_{\ell^*} | c_{\ell^*} \rangle - \mathcal{D}(\mathcal{B}_{\delta\lambda_1^*}(S) \cap [-q/2, q/2]^m) \\ \leq O\left(\frac{\sqrt{m}s^{m/2}(V \text{Lip}(\mathbf{f}) + s^2)}{q}\right) + o(\eta) \end{aligned} \quad (3.43)$$

As the only non-negligible error is caused by Equation (3.33), provided that $\delta V \lambda_1^* / s \geq \sqrt{m + \log(\eta^{-1})}$, $s \geq 4\sqrt{m \log m + \log(\eta^{-1})}$ and $q/s \geq \sqrt{m + \log(\eta^{-1})}$.

Remark 3.18. Note that we chose for $S_{\text{fin}} = S \cap [-q/4, q/4]^m$ the box $[-q/4, q/4]^m$, whereas in the analysis we used the box $[-q/2, q/2]^m$. This is to crudely include also all points that are $\delta\lambda_1^*$ -close to dual lattice vectors.

Final theorem

Assembling all errors, we obtain the following theorem.

Theorem 3.7. *Algorithm 2 solves the Dual Lattice Sampling Problem with parameters η and δ ; it uses one call to the Gaussian superposition subroutine (see Theorem 3.12), one quantum oracle call to \mathbf{f} , $mQ + n$ qubits, and $O(mQ \log(mQ))$ quantum gates, where*

$$Q = O(m \log(m)) + O\left(\log\left(\frac{a}{\eta \cdot \delta \lambda_1^*}\right)\right). \quad (3.27)$$

Proof. In Algorithm 2, two quantum registers are used: one to encode the grid \mathbb{D}^m and another one for the storage of the state of the continuous hidden subgroup oracle $|\mathbf{f}(x)\rangle$. As the grid has q^m points, we need $m \log q$ qubits to encode it. For the oracle state it is assumed that it can be stored in n qubits, thus arriving at a total of $mQ + n$ qubits, where $Q = \log q$. Apart from constructing the initial Gaussian superposition, the only part of Algorithm 2 that uses quantum gates is the quantum Fourier transform on the grid register consisting of mQ qubits. Using a result of Hallgren et al., a sufficient approximation of this quantum Fourier transform can be obtained using only $O(mQ \log(mQ))$ elementary quantum gates [HH00].

To compute the value of $Q = \log(q)$, we instantiate the parameters $s = 4\sqrt{m \log m + \log(\eta^{-1})}$ and $V = \frac{4}{\delta \lambda_1^*} \cdot (m \log m + \log(\eta^{-1}))$. This implies $s \geq 4\sqrt{m \log m + \log(\eta^{-1})}$ and $\delta V \lambda_1^*/s \geq \sqrt{m + \log(\eta^{-1})}$, making the errors from Equations (3.31), (3.32), (3.39) and (3.41) all negligible compared to η . To get the errors from Equation (3.33) and Equation (3.43) well below η , we put

$$\log q = Q = O\left(m \log(s) + \log\left(\frac{V \text{Lip}(\mathbf{f})}{\eta}\right)\right).$$

Writing out the instantiations of s and V and grouping the resulting expressions properly, we arrive at Equation (3.27). Here we use the fact that, for all $\eta > 0$ and $m \in \mathbb{N}$, $m \left(\log(m \log m + \log(1/\eta)) \right) \in O(m \log m + \log(1/\eta))$.

□

3.6. From Sampling to Full Dual Lattice Recovery

We have so far focused on approximate sampling dual lattice points with probability weights $\langle c_{\ell^*} | c_{\ell^*} \rangle$ for $\ell^* \in \Lambda^*$, regardless of how useful this distribution may be. Indeed, until now, it could be that the function $\mathbf{f} : \mathbb{R}^m / \Lambda \rightarrow \mathcal{S}$ is constant, and therefore that all weight is concentrated on $0 \in \Lambda^*$. We would like now make sure we can reconstruct (approximately) Λ^* from such samples, i.e., that a sufficient number of sampled vectors from Λ^* will generate it. Informally, an equivalent condition is that the weight $\langle c_{\ell^*} | c_{\ell^*} \rangle$ is not concentrated on any proper sublattice $M^* \subsetneq \Lambda^*$. This is exactly what happens if the oracle function \mathbf{f} is separating, i.e., is not too constant.

More formally, we give the following sufficient conditions for a distribution to be useful as a (approximate) lattice sampling distribution.

Definition 3.19. Let $L \subseteq \mathbb{R}^m$ be a full-rank lattice. A distribution \mathcal{D} on L is called p -evenly distributed whenever $\Pr_{v \leftarrow \mathcal{D}}[v \in L'] \leq p$ for any proper sublattice $L' \subsetneq L$.

Definition 3.20. Let $L \subseteq \mathbb{R}^m$ be a full-rank lattice. A distribution \mathcal{D} on L is called (R, q) -concentrated whenever $\Pr_{v \leftarrow \mathcal{D}}[\|v\| \geq R] \leq q$.

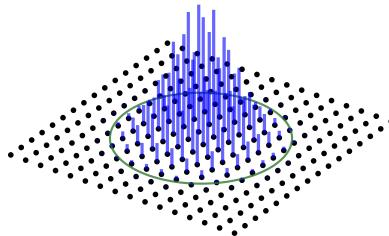


Figure 3.7.: An example of a (R, q) -concentrated distribution, where R is the radius of the green circle and $q = 0.05$, i.e., less than 5 percent of the weight lies outside the circle. Note that this Gaussian distribution is also 0.5-evenly distributed.

The following lemma states that an evenly distributed and well-concentrated distribution on a lattice L should eventually output a full generating set of

3. The Continuous Hidden Subgroup Problem

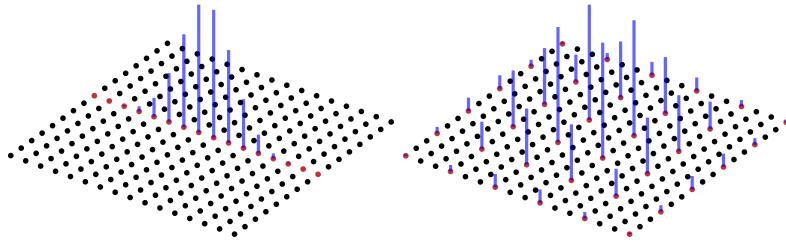


Figure 3.8.: Both these distributions are not p -evenly distributed for any $p < 1$, as the strict sublattices indicated by the red points have all of the weight.

that lattice, and gives a precise probabilistic upper bound on the number of samples needed.

Lemma 3.21. *Let $L \subseteq \mathbb{R}^m$ be a full-rank lattice with a p -evenly distributed and (R, q) -concentrated distribution \mathcal{D} with $R \geq \det(L)^{1/m}$. Denote by S the random variable defined by the number of samples that needs to be drawn from \mathcal{D} such that the samples together generate L as a lattice. Then, for all $\alpha > 0$,*

$$\Pr \left[S > (2 + \alpha) \cdot \frac{(t + m)}{1 - p - q} \right] \leq \exp(-\alpha(t + m)/2)$$

where $t = m \log_2(R) - \log_2(\det(L)) \geq 0$.

Proof. First, we define the following sublattices of L , for any $v_1, \dots, v_{j-1} \in L$.

$$L_{v_1, \dots, v_{j-1}} = \begin{cases} \text{span}_{\mathbb{R}}(v_1, \dots, v_{j-1}) \cap L & \text{if } \dim(\text{span}_{\mathbb{R}}(v_1, \dots, v_{j-1})) < m \\ \mathbb{Z}v_1 + \dots + \mathbb{Z}v_{j-1} & \text{otherwise.} \end{cases}$$

Consider a sequence of samples $(v_i)_{i>0}$ (from \mathcal{D}). We call v_j ‘good’ whenever $\|v_j\| \leq R$ and $v_j \notin L_{v_1, \dots, v_{j-1}}$. We argue that we need at most $m + t$ good vectors to generate L .

Denote L' for the lattice generated by the $m + t$ good vectors. Then the first m good vectors ensure that L' is of rank m , whereas the last t good vectors will reduce the index of the L' lattice in L . Calculating determinants,

using the fact that all good vectors are bounded by R , we have $\det(L') \leq R^m/2^t \leq \det(L)$. This yields $L' = L$.

Denote by X the random variable having the negative binomial distribution with success probability $p + q$ and number of ‘failures’ $m + t$. That is, X is the number of independent samples from a $(p + q)$ -Bernoulli distribution until $m + t$ ‘failures’¹ are obtained. We argue that the random variable S is dominated by the random variable X , i.e., $\Pr[S > x] \leq \Pr[X > x]$ for every $x \in \mathbb{N}$.

Again, consider a sequence of samples $(v_i)_{i>0}$ (from \mathcal{D}). The probability of v_j being a ‘good’ vector is at least $1 - p - q$, by the fact that \mathcal{D} is (R, q) -concentrated and p -evenly distributed. Because at most $m + t$ ‘good’ vectors are needed to generate the whole lattice, S is indeed dominated by X . Therefore, for any $k \in \mathbb{N}$,

$$\begin{aligned} \Pr \left[S > \frac{t + m + k}{1 - p - q} \right] &\leq \Pr \left[X > \frac{t + m + k}{1 - p - q} \right] \leq \Pr [B < m + t] \\ &\leq \exp \left(-\frac{1}{2} \frac{k^2}{t + m + k} \right) \end{aligned} \tag{3.44}$$

where B is binomially distributed with $\lfloor \frac{t+m+k}{1-p-q} \rfloor$ trials and success probability $1 - p - q$. The first inequality follows from the fact that S is upper bounded by X . The second inequality comes from the close relationship between the negative binomial distribution and the binomial distribution [GKP94, Ch. 8, Example 17]. The last inequality follows from the Chernoff bound. Putting $k = (1 + \alpha)(t + m)$ into Equation (3.44) yields the claim. \square

We conclude this section by relating the parameters (a, r, ϵ) of the HSP oracle (Definition 3.2) $\mathbf{f} : \mathbb{R}^m/\Lambda \rightarrow \mathcal{S}$ to how equally-distributed and well-concentrated the distribution \mathcal{D}_{ideal} on Λ^* is, arising from the Fourier coefficients of the oracle function \mathbf{f} . The exact relation is stated in Proposition 3.24, but we first need two technical lemmas to help us proving this relation.

¹In our case, the failures are the ‘good’ vectors. We nonetheless chose the word ‘failure’ because it is standard nomenclature for the negative binomial distribution.

3. The Continuous Hidden Subgroup Problem

Lemma 3.22. *Let Λ be a lattice, and let $M \supsetneq \Lambda$ a proper super-lattice of Λ . Then there exists a $v \in M$ such that $d(v, \Lambda) \geq \lambda_1(\Lambda)/3$.*

Proof. Let $w \in M$ be the shortest non-zero vector in M and write $\|w\| = \alpha \lambda_1(\Lambda)$ for $\alpha \leq 1$. We consider two cases depending on the value of $\alpha \in (0, 1]$. If $\alpha \geq 1/3$, choose an element $v \in M \setminus \Lambda$ arbitrarily. This element satisfies $d(v, \Lambda) \geq \lambda_1(\Lambda)/3$, since

$$\begin{aligned} d(v, \Lambda) &= d(v + \Lambda, 0) = d((v + \Lambda) \setminus 0, 0) \\ &\geq d(M \setminus 0, 0) = \alpha \cdot \lambda_1(\Lambda) \geq \lambda_1(\Lambda)/3. \end{aligned}$$

If, on the other hand, $\alpha < 1/3$, then $v = \lceil \frac{1}{3\alpha} \rceil \cdot w \in M$ satisfies $d(v, \Lambda) \geq \lambda_1(\Lambda)/3$. One can deduce this by observing that

$$\|v\| = \lceil \frac{1}{3\alpha} \rceil \cdot \alpha \cdot \lambda_1(\Lambda) \in [\frac{1}{3} \cdot \lambda_1(\Lambda), \frac{2}{3} \cdot \lambda_1(\Lambda)],$$

which in particular implies that $\|v - \ell\| \geq \|\ell\| - \|v\| \geq \frac{1}{3} \cdot \lambda_1(\Lambda)$, for all $\ell \in \Lambda$, i.e., $d(v, \Lambda) \geq \frac{1}{3} \lambda_1(\Lambda)$. \square

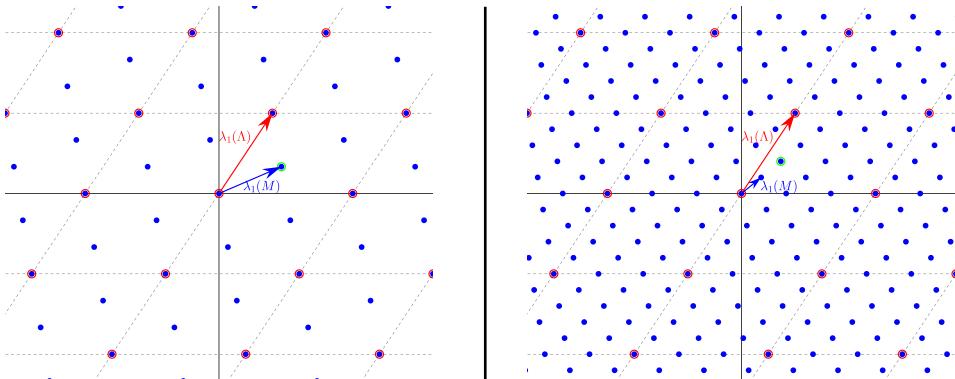


Figure 3.9.: The two cases of Lemma 3.22 are depicted here, where the smaller lattice Λ consists of the points inside the red circles. The blue super-lattice M satisfies $\lambda_1(M) = \alpha \lambda_1(\Lambda)$ for some $\alpha > 1/3$ in the left picture and for some $\alpha < 1/3$ in the right picture. In both cases, an element $v \in M$ for which holds $d(v, \Lambda) > \frac{1}{3} \cdot \lambda_1(\Lambda)$ can be reasonably found. Examples of such $v \in M$ are marked with a green circle.

Lemma 3.23. *Let Λ be a lattice and $M \supsetneq \Lambda$ a proper super-lattice of Λ . Then the number $N = \left| \left\{ c \in M/\Lambda \mid d(c, \Lambda) < \frac{1}{6}\lambda_1(\Lambda) \right\} \right|$ of close cosets is at most $\frac{1}{2} \cdot |M/\Lambda|$.*

Proof. By Lemma 3.22 there exists a $v \in M$ such that $d(v, \Lambda) \geq \frac{1}{3}\lambda_1(\Lambda)$. Denoting $T = \left\{ c \in M/\Lambda \mid d(c, \Lambda) < \frac{1}{6}\lambda_1(\Lambda) \right\}$, we can deduce that $T \cup (T+v)$ is a disjoint union in M/Λ . Indeed, elements $c \in T$ satisfy $d(c, \Lambda) < \frac{1}{6}\lambda_1(\Lambda)$, whereas $c' \in T+v$ satisfy $d(c', \Lambda) \geq d(v, \Lambda) - \frac{1}{6}\lambda_1(\Lambda) \geq \frac{1}{6}\lambda_1(\Lambda)$. Therefore $N = |T| \leq \frac{1}{2}|M/\Lambda|$. \square

Proposition 3.24. *Let $\mathbf{f} : \mathbb{R}^m \rightarrow \mathcal{S}$ be an (a, r, ϵ) -HSP oracle of the full-rank lattice $\Lambda \subset \mathbb{R}^m$, with $r \leq \lambda_1(\Lambda)/6$. Let $\mathcal{D}_{\mathbf{f}}$ be the distribution supported by Λ^* , with weight $\langle c_{\ell^*} | c_{\ell^*} \rangle$ at $\ell^* \in \Lambda^*$, where $|c_{\ell^*}\rangle$ are the vectorial Fourier coefficients of the function \mathbf{f} . Then $\mathcal{D}_{\mathbf{f}}$ is both $(\frac{1}{2} + \epsilon)$ -evenly distributed and $(R, \frac{ma^2}{4\pi^2 R^2})$ -concentrated for any $R > 0$.*

Proof. The distribution $\mathcal{D}_{\mathbf{f}}$ being $(R, \frac{ma^2}{4\pi^2 R^2})$ -concentrated for any $R > 0$ is a direct consequence of Corollary 2.34. For the $(\frac{1}{2} + \epsilon)$ -evenly distributed part, we argue as follows. Let M^* be any strict sublattice of Λ^* , and let M be its dual, which is then a superlattice of Λ . Put $\mathbf{f}|^{\mathbb{R}^m/M}(x) = \frac{1}{|M/\Lambda|} \sum_{v \in M/\Lambda} \mathbf{f}(x + v)$, the periodization of \mathbf{f} with respect to \mathbb{R}^m/M (c.f. Definition 2.3). We have the following sequence of equalities, of which the second follows from the Poisson summation formula (see Theorem 2.4) and

3. The Continuous Hidden Subgroup Problem

the third from Parseval's theorem (see Equation (2.5)).

$$\begin{aligned}
& \sum_{v^* \in M^*} \langle c_{v^*} | c_{v^*} \rangle \\
&= \left\| \mathcal{F}_{\mathbb{R}^m / \Lambda} \{ \mathbf{f} \} \right\|_{M^*}^2 = \left\| \mathcal{F}_{\mathbb{R}^m / M} \{ \mathbf{f} |_{\mathbb{R}^m / M} \} \right\|_{M^*}^2 \\
&= \|\mathbf{f} |_{\mathbb{R}^m / M}\|_{\mathbb{R}^m / M} = \frac{1}{\det M} \int_{x \in \mathbb{R}^m / M} \langle \mathbf{f} |_{\mathbb{R}^m / M} | \mathbf{f} |_{\mathbb{R}^m / M} \rangle dx, \\
&= \frac{1}{|M/\Lambda|^2} \sum_{v, w \in M/\Lambda} \underbrace{\frac{1}{\det M} \int_{x \in \mathbb{R}^m / M} \langle \mathbf{f}(x+v) | \mathbf{f}(x+w) \rangle dx}_{I_{v,w}} \\
&= \frac{1}{|M/\Lambda|^2} \sum_{\substack{v, w \in M/\Lambda \\ d_{\mathbb{R}^m / \Lambda}(v, w) < r}} I_{v,w} + \frac{1}{|M/\Lambda|^2} \sum_{\substack{v, w \in M/\Lambda \\ d_{\mathbb{R}^m / \Lambda}(v, w) \geq r}} I_{v,w}. \quad (3.45)
\end{aligned}$$

By the definition of an (a, r, ϵ) -oracle, we have that $|I_{v,w}| \leq \epsilon$ whenever $d_{\mathbb{R}^m / \Lambda}(v, w) \geq r$. In the rest of the cases we have $|I_{v,w}| \leq 1$, because \mathbf{f} maps to the unit sphere. Equation (3.45) is therefore bounded by $\frac{|M/\Lambda \cap \mathcal{B}_r|}{|M/\Lambda|} + \epsilon$, where \mathcal{B}_r is the open unit ball around zero with radius r . By Lemma 3.23, we have $\frac{|M/\Lambda \cap r\mathcal{B}|}{|M/\Lambda|} \leq \frac{1}{2}$ for $r \leq \lambda_1(\Lambda)/6$. Summarizing, we derive

$$\sum_{v^* \in M^*} \langle c_{v^*} | c_{v^*} \rangle \leq \frac{1}{2} + \epsilon.$$

Since M^* was chosen arbitrarily, we can conclude that $\mathcal{D}_{\mathbf{f}}$ is $(\frac{1}{2} + \epsilon)$ -evenly distributed. \square

Remark 3.25. A similar reasoning happens in [Reg04a, Lecture 12], though it specifically targets the discrete Gaussian distribution on lattices. Despite being not general enough for our purposes, it may well be helpful for optimizing a future specialization.

Theorem 3.9. Let $\mathbf{f} : \mathbb{R}^m \rightarrow \mathcal{S}$ be an (a, r, ϵ) -HSP oracle of the full-rank lattice $\Lambda \subset \mathbb{R}^m$, with $r \leq \lambda_1(\Lambda)/6$ and $\epsilon < 1/4$. Let $\mathcal{D}_{\mathbf{f}}$ be the distribution supported by Λ^* , with weight $\langle c_{\ell^*} | c_{\ell^*} \rangle$ at $\ell^* \in \Lambda^*$, where $|c_{\ell^*}\rangle$ are the vectorial

Fourier coefficients of the function \mathbf{f} .

Then, with overwhelming probability, we need at most

$$O\left(m \log_2 (ma \cdot \det(\Lambda)^{1/m})\right)$$

samples from \mathcal{D}_f to fully generate the lattice Λ^* .

Proof. Apply Proposition 3.24 with $R = \sqrt{m} \cdot \text{Lip}(\mathbf{f})$ to deduce that \mathcal{D}_f is $3/4$ -evenly distributed and $(\sqrt{m} \text{Lip}(\mathbf{f}), 1/(4\pi^2))$ -concentrated. Subsequently, we apply Lemma 3.21 with² $p = 3/4$, $q = 1/(4\pi^2)$, $R = \sqrt{m} \cdot \text{Lip}(\mathbf{f})$ and $t = m \log_2(\sqrt{m} \text{Lip}(\mathbf{f})) - \log_2(\det(\Lambda^*))$, to obtain

$$\Pr [S > (2 + \alpha) \cdot 5 \cdot (t + m)] \leq \exp(-\alpha(t + m)/2).$$

Writing out t (which is larger than 0), noticing that $\text{Lip}(\mathbf{f}) \leq a$, and absorbing m into the big-O, we obtain the result with exponentially small error probability. \square

3.7. Recovering a Basis of the Primal Lattice

The last problem that needs to be resolved is how to obtain an approximate basis \tilde{B} of the primal lattice Λ , given a set of approximate generators \tilde{G} of the dual lattice Λ^* . Also, we would like to know how the approximation errors of \tilde{G} and \tilde{B} are related.

Recovering the approximate basis \tilde{B} proceeds by two steps. The first step consists of applying the Buchmann-Pohst algorithm [BP89] twice to the set of generators \tilde{G} , yielding an approximate basis \tilde{D} of the dual lattice Λ^* whose errors are relatively easy to analyze. The second step consists of inverting and transposing the square matrix \tilde{D} . This yields an approximate basis \tilde{B} for the primal lattice Λ .

²In order to apply Lemma 3.21, we need to verify that $R = \sqrt{m} \text{Lip}(\mathbf{f}) \geq \det(\Lambda^*)^{1/m}$. By Remark 3.5, we have $\sqrt{m} \text{Lip}(\mathbf{f}) \geq \frac{\sqrt{m}(1-\epsilon)}{r} \geq \frac{3\sqrt{m}}{\lambda_1(\Lambda)} \geq 3 \det(\Lambda)^{-1/m} = 3 \det(\Lambda^*)^{1/m}$, where we use $r \leq \lambda_1(\Lambda)/6$ and Minkowski's inequality.

3. The Continuous Hidden Subgroup Problem

The next two subsections follow above summary, and consist of theorems that indicate the decline in precision after each step.

In this particular section, we use row notation for matrices, i.e., any row represents a vector. The matrix of generators \tilde{G} is an $k \times m$ matrix, thus consisting of k generators. We assume that the lattice Λ (and thus Λ^* as well) is of full rank m , meaning that $k \geq m$ and that the resulting bases \tilde{D} and \tilde{B} must be $m \times m$ square matrices. We denote by $\|M\|_\infty$ the matrix norm induced by the infinity norm, explicitly defined as

$$\|M\|_\infty := \max_{1 \leq i \leq m} \sum_{j=1}^n |m_{ij}|.$$

3.7.1. An Approximate Well-conditioned Basis of the Dual

Obtaining an approximate and well-conditioned basis of the dual proceeds by means of the Buchmann-Pohst algorithm, which is rigorously analyzed by Buchmann and Kessler [BK96, Sec. 4]. This algorithm consists of concatenating the generating matrix by a scaled identity matrix and applying the LLL lattice reduction algorithm. As described after the proof of [BK96, Thm. 4], this particular algorithm is actually applied twice, once on the matrix of generators \tilde{G} and once again on the resulting intermediate approximate basis \tilde{D} to achieve a new basis whose errors are easier to analyze. From now on, we will refer to applying this procedure twice as the Buchmann-Pohst algorithm. From [BK96] we can extract the following result.

Theorem 3.26. *Let $\tilde{G} = G + \Delta_G$ be an approximation of a $k \times m$ matrix of generators G of the full-rank lattice Λ^* , with $\|\Delta_G\|_\infty < \gamma < \frac{\lambda_1^* \cdot \det(\Lambda^*)}{2^{O(km)} \cdot \|\tilde{G}^*\|_\infty^m}$. Then, on input³ $[\tilde{G} \mid \gamma \cdot I]$, the Buchmann-Pohst algorithm outputs an LLL-reduced matrix $[\tilde{D} \mid \gamma \cdot M]$, with $\tilde{D} = D + \Delta_D$ being an approximate basis of Λ^* , where both $\|\Delta_D\|_\infty$ and $\|\gamma \cdot M\|_\infty$ are upper bounded by*

$$\frac{2^{O(km)} \|\tilde{G}^*\|_\infty^{m+1}}{\lambda_1^* \cdot \det(\Lambda^*)} \cdot \gamma$$

³Here, the I is the $k \times k$ identity matrix.

Proof. As already mentioned, applying the Buchmann-Pohst algorithm on \tilde{G} takes two reduction steps. The first reduction step is applied on $[\tilde{G} | \gamma \cdot I]$ and yields an intermediate basis $\tilde{D}_0 = M_0 \tilde{G}$ and the second step is applied on $[\tilde{D}_0 | \gamma I]$ and yields the final basis $\tilde{D} = M \tilde{D}_0 = M M_0 \tilde{G}$. Here, M and M_0 are unimodular matrices.

The fact that the matrix $[\tilde{D} | \gamma \cdot M]$ is the output of the second step, proves that this must be an LLL-reduced basis (note that $\gamma \cdot M$ is just the matrix M scaled by the scalar γ). From [BK96, Cor. 4.1], we deduce that both $\|M\|_\infty$ and $\|MM_0\|_\infty$ are bounded by $2^{k-1}(\sqrt{mk}+2) \cdot \lambda' \alpha' / \lambda_1(\Lambda^*)$, given that⁴ $\gamma < \frac{\lambda_1(\Lambda^*) \det(\Lambda^*)}{(\sqrt{mk}+2) \cdot \lambda' \cdot 2^{\frac{k-3}{2}}}$. Putting in the actual values of $\alpha' = (\sqrt{mk}+2)2^{\frac{k-1}{2}} \cdot \|\tilde{G}\|_\infty$ and

$$\lambda' = \lambda(\sqrt{mk}+2)^m 2^{\frac{k-1}{2}m} = (k\sqrt{m}/2 + \sqrt{k})(\sqrt{mk}+2)^m 2^{\frac{k-1}{2}m} \frac{\|\tilde{G}\|_\infty^m}{\det(\Lambda^*)}$$

yields the bound on $\|\gamma M\|_\infty$ and the assumption on γ . For the bound on Δ_D , notice that $\|\Delta_D\|_\infty = \|MM_0\Delta_G\|_\infty \leq \|MM_0\|_\infty \|\Delta_G\|_\infty$ and by assumption $\|\Delta_G\|_\infty \leq \gamma$. \square

For small enough γ , the LLL-reduced basis $[\tilde{D} | \gamma \cdot M]$ is very close to $[D | 0]$. One of the main results of Chang, Stehlé and Villard [CSV12, Cor. 5.7] states that the close matrix $[D | 0]$ must then also be ‘weakly LLL-reduced’. This knowledge can then be used to show that this basis D is well-conditioned.

Lemma 3.27. *Let $[\tilde{D} | \gamma M] = [D | 0] + [\Delta_D | \gamma M]$ be an LLL-reduced basis with $\|[\Delta_D | \gamma M]\|_\infty \leq \mu \cdot (3/\sqrt{2})^{-3m} \|\tilde{D}\|_\infty$ for some $\mu < 1$. Then D is (d, η', θ') -weakly LLL-reduced as in [CSV12, Def. 5.1], with $d = \frac{3}{4} + O(2^{-m}\mu)$, $\eta = \frac{1}{2} + O(2^{-m}\mu)$ and $\theta = O(2^{-m}\mu)$.*

Corollary 3.28. *Let $[\tilde{D} | \gamma M] = [D | 0] + [\Delta_D | \gamma M]$ be an LLL-reduced basis with $\|[\Delta_D | \gamma M]\|_\infty \leq \mu \cdot (3/\sqrt{2})^{-3m} \|\tilde{D}\|_\infty$ for some $\mu < 1$ (i.e.,*

⁴See [BK96, Thm. 4.1], where λ needs to be replaced by λ' , as described in the text after the proof of [BK96, Thm. 4.2]). These variables λ and λ' are from [BK96, Prop. 3.2], and not directly related to the minima of the lattices involved.

3. The Continuous Hidden Subgroup Problem

satisfies the same assumptions as in Lemma 3.27). Then

$$\|D^{-1}\|_{\infty} \leq \frac{8^m}{\lambda_1(\Lambda^*)}.$$

Proof. We can decompose $D = RVQ$, with Q orthonormal, V diagonal with diagonal entries $\|d_i^*\|$ and R lower triangular with ones on the diagonal. Here, d_i^* are the Gram-Schmidt orthogonalized basis vectors of D .

By the fact that the matrix norm is submultiplicative, we have

$$\|D^{-1}\|_{\infty} \leq \|R^{-1}\|_{\infty} \|V^{-1}\|_{\infty} \|Q^{-1}\|_{\infty} = \|R^{-1}\|_{\infty} \|V^{-1}\|_{\infty} \leq \frac{\|R^{-1}\|_{\infty}}{\min_i \|d_i^*\|}.$$

By Lemma 3.27, D is weakly (d, η, θ) -LLL-reduced with $d = \frac{3}{4} + O(2^{-m}\mu)$, $\eta = \frac{1}{2} + O(2^{-m}\mu)$ and $\theta = O(2^{-m}\mu)$. Therefore, by [CSV12, Thm. 5.4], taking $\alpha = 2 > \sqrt{2}$ for simplicity, we know that $\lambda_1(\Lambda^*) \leq \|d_1\| \leq 2^m \min_i \|d_i^*\|$, so that $1/\min_i \|d_i^*\| \leq 2^m \lambda_1(\Lambda^*)^{-1}$. In the end of the proof of [CSV12, Lm. 5.5], we see⁵ that

$$\|R^{-1}\|_{\infty} \leq \frac{(1+\alpha)(1+\eta+\theta)^m \alpha^m}{(1+\eta+\theta)\alpha - 1} \leq 4^m,$$

by taking $\alpha = 2$, $\eta = 1/2 + O(2^{-m}\mu)$ and $\theta = O(2^{-m}\mu)$. This yields the claim. \square

3.7.2. Inverting the Dual Approximate Basis

As the basis \tilde{D} constructed in the previous subsection is a basis of the *dual* lattice Λ^* , we need to invert and transpose it to get an approximate basis of the primal lattice Λ . In other words, the basis that we would like to approximate is $B = D^{-\top}$, by means of computing $\tilde{B} = \tilde{D}^{-\top}$. Though, inverting an approximate matrix induces errors closely related with the matrix norm of the inverse of the exact basis. More precisely, we have the following result [BKK17, Cor. 7.2, Eq. (7.46)]

⁵In [CSV12, Lm. 5.5], the unit-diagonal lower triangular matrix is denoted \bar{R} , and the bound is about \bar{R}^{-1}

Theorem 3.29. Let $\tilde{D} = D + \Delta_D$ with $\|\Delta_D\|_\infty \cdot \|D^{-1}\|_\infty < \frac{1}{2}$, and denote $B = D^{-\top}$ and $\tilde{B} = \tilde{D}^{-\top}$ (where $D^{-\top}$ is the inverse transpose of D). Then we have

$$\|B - \tilde{B}\|_\infty \leq 2\|D^{-1}\|^2\|\Delta_D\|_\infty.$$

3.7.3. Combining the Errors and Tuning the Parameters

Theorem 3.10. There exists a polynomial time algorithm, that, for any matrix $G \in \mathbb{R}^{k \times m}$ of k generators of a (dual) lattice Λ^* , and given an approximation $\tilde{G} = G + \Delta_G \in \mathbb{Q}^{k \times n}$, computes an approximation $\tilde{B} = B + \Delta_B$ of a basis B of the primal lattice Λ , such that

$$\|\Delta_B\|_\infty \leq \frac{2^{O(mk)} \cdot \|\tilde{G}\|_\infty^{m+1}}{(\lambda_1^*)^3 \cdot \det(\Lambda^*)} \cdot \|\Delta_G\|_\infty,$$

under the assumption that $\|\Delta_G\|_\infty < \frac{\min(1, (\lambda_1^*)^2) \cdot \det(\Lambda^*)}{2^{O(km)} \cdot \|\tilde{G}\|_\infty^{m+1}}$.

Proof. For the moment, assume that the full output⁶ $[\tilde{D} \mid \gamma M] = [D \mid 0] + [\Delta_D \mid \gamma M]$ of the Buchmann-Pohst algorithm satisfies $\|[\Delta_D \mid \gamma M]\|_\infty \leq \mu(3/\sqrt{2})^{-3m}\|\tilde{D}\|_\infty$ for some $\mu < 1$ and $\|\Delta_D\|_\infty\|D^{-1}\|_\infty < 1/2$. Then, by applying Theorem 3.29, Corollary 3.28 and Theorem 3.26 subsequently, we obtain

$$\|\Delta_B\|_\infty \leq 2\|D^{-1}\|_\infty^2\|\Delta_D\|_\infty \leq \frac{2^{6m+1}}{(\lambda_1^*)^2}\|\Delta_D\|_\infty \leq \frac{2^{O(mk)} \cdot \|\tilde{G}\|_\infty^{m+1}}{(\lambda_1^*)^3 \cdot \det(\Lambda^*)} \cdot \gamma.$$

It remains to prove that assumptions in the beginning of this proof are indeed fulfilled. By Theorem 3.26, we have

$$\|[\Delta_D \mid \gamma M]\|_\infty \leq \frac{2^{O(mk)} \cdot \|\tilde{G}\|_\infty^{m+1}}{\lambda_1^* \cdot \det(\Lambda^*)} \cdot \gamma < O(1),$$

⁶In reality, the Buchmann-Pohst algorithm is applied with the largest precision such that all required assumptions hold. So the costs of applying the LLL-algorithm does not involve the precision $\|\Delta_G\|_\infty$.

3. The Continuous Hidden Subgroup Problem

and by Theorem 3.29, we have

$$\|\Delta_D\|_\infty \|D^{-1}\|_\infty \leq \|\Delta_D\|_\infty \frac{2^{3m}}{\lambda_1^*} \leq \frac{2^{O(mk)} \cdot \|\tilde{G}\|_\infty^{m+1}}{(\lambda_1^*)^2 \cdot \det(\Lambda^*)} \cdot \gamma < O(1).$$

So choosing γ appropriately small, the assumptions of Theorem 3.29, Corollary 3.28 and Theorem 3.26 are all fulfilled. \square

4. Random Walks on Arakelov Ray Class Groups

4.1. Summary

The Arakelov class group.

The Arakelov class group (denoted Pic_K^0) is a combination of the unit torus $T = \text{Log } K_{\mathbb{R}}^0 / \text{Log}(\mathcal{O}_K^\times)$ and the class group Cl_K . The exponent 0 in $K_{\mathbb{R}}^0$ refers to elements of algebraic norm 1 (i.e., modulo renormalization), while the subscript \mathbb{R} indicates that we are working in the tensor product of K with \mathbb{R} over \mathbb{Q} . By ‘a combination’ we mean that there is a short exact sequence

$$0 \longrightarrow T \longrightarrow \text{Pic}_K^0 \longrightarrow \text{Cl}_K \longrightarrow 0.$$

That is, T is (isomorphic to) a subgroup of Pic_K^0 , and Cl_K is isomorphic to the quotient Pic_K^0 / T . Summarizing, the Arakelov class group is an abelian group which combines an uncountable but compact part T and a finite part Cl_K ; topologically, it should be thought of as $|\text{Cl}_K|$ many disconnected copies of the torus T .

The Arakelov ray class group.

In this chapter we actually consider a more general group, an Arakelov analogue of the finite *ideal ray class group* $\text{Cl}_K^m = \mathcal{I}_K^m / \text{Princ}_K^m$; which is the

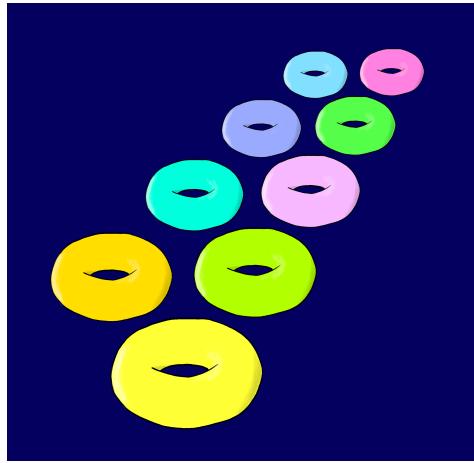


Figure 4.1.: The Arakelov class group can be thought of as $|\text{Cl}_K|$ copies of the logarithmic unit torus.

reason why it is named the Arakelov *ray* class group. It is defined likewise via an exact sequence,

$$0 \longrightarrow T^{\mathfrak{m}} \longrightarrow \text{Pic}_{K^{\mathfrak{m}}}^0 \longrightarrow \text{Cl}_K^{\mathfrak{m}} \longrightarrow 0,$$

where $T^{\mathfrak{m}} = \text{Log } K_{\mathbb{R}}^0 / \text{Log}(\mathcal{O}_K^{\times} \cap K^{\mathfrak{m},1})$. Here, $K^{\mathfrak{m},1}$ is called the *ray* of \mathfrak{m} , the multiplicative subgroup of K^* generated by elements that are congruent to 1 modulo \mathfrak{m} . This Arakelov ray class group has essentially the same structure as the ‘normal’ Arakelov class group (which can be recovered by taking $\mathfrak{m} = \mathcal{O}_K$), in the sense that it can also be thought of as $|\text{Cl}_K^{\mathfrak{m}}|$ many disconnected copies of the (larger) torus $T^{\mathfrak{m}}$.

Random walks on the Arakelov ray class group.

In this chapter we study the process of a random walk on this Arakelov ray class group, which can be described best by using the correspondence with Arakelov ray class group elements with *ideal lattices*. This random walk process consists of multiplying the input ideal lattice by a certain number of random prime ideals of bounded norm, followed by a slight disturbance of

the geometry of the ideal lattice. The end distribution over $\text{Pic}_{K^m}^0$ is called the *random walk distribution*.

The main question to be solved is for which choice of parameters this random walk distribution is close to the uniform distribution. These parameters involve the maximum norm of the prime ideals, the number of prime ideals to multiply with and the magnitude of the geometrical disturbance.

Fourier analysis on the Arakelov ray class group.

Because the Arakelov ray class group is abelian and compact, this question is tackled by resorting to Fourier analysis: uniformity is demonstrated by showing that all the Fourier coefficients of the distribution resulting from the random walk tend to 0 except for the coefficient associated with the trivial character.

This argument can be roughly described as follows. The act of multiplying by random prime ideals can be described by a so-called Hecke operator, a linear Hermitian operator that has the characters of $\text{Pic}_{K^m}^0$ as eigenfunctions. Assuming the Extended Riemann Hypothesis, one can show that all eigenvalues of the non-trivial eigenfunctions are bounded sufficiently below one, except for specific ‘high-frequency eigenfunctions’. Choosing an initial distribution (the ‘geometrical disturbance’) that lacks those high-frequency eigenfunctions, e.g., a Gaussian, applying the Hecke operator sufficiently many times yields a near-uniform distribution (see Figure 4.7).

The preciseness of this argument allows us to very tightly estimate bounds for all parameters involved in order to achieve a nearly uniformly random final distribution.

4.2. Introduction

One of the more important concepts that occurs in complexity theory, number theory and modern cryptography is that of a *structured lattice*. The

most elementary examples of such structured lattices are *ideal lattices*, which are derived from ideals of a number field. It is a well-known fact among number theorists that the set of all ideal lattices up to K -linear isometry (or, equivalently, Hermitian line bundles) associated with a fixed number field K forms a compact abelian group, the *Arakelov class group* Pic_K^0 (e.g. [Sch08]). The motivation for studying this particular Arakelov class group in this thesis came from two directions.

The first direction relates to number theory and involves a computational problem related to the density of prime ideals. Namely: for a given ideal \mathfrak{a} of \mathcal{O}_K , find an element $\alpha \in \mathfrak{a}$ with a predetermined bounded length such that $(\alpha)/\mathfrak{a}$ is a prime ideal. The most straightforward way to obtain such an element without class group computations is by *randomly* sampling an element α in the intersection of the ideal \mathfrak{a} with a large box, and just simply *hope* that the ideal $(\alpha)/\mathfrak{a}$ is prime. For a fixed ideal \mathfrak{a} we cannot prove that this approach is efficient, but for an ideal lattice \mathfrak{a} that is *uniformly randomly distributed* in the Arakelov class group Pic_K^0 we can actually reasonably lower bound the success probability of this approach (see Chapter 6). The remaining question is: can we efficiently transform a fixed ideal \mathfrak{a} into a random ideal in the Arakelov class group without changing too much of its properties?

The second direction relates more to complexity theory and cryptography, and involves the computational hardness of the *Hermite Shortest Vector Problem* on ideal lattices of a fixed number field. An interesting question here is, for a fixed number field K , whether there are ideal lattices that are significantly harder to solve Hermite-SVP in than for other ideal lattices. A natural way to approach this question is by comparing the hardness of Hermite-SVP on a fixed ideal lattice with the hardness of Hermite-SVP on an *average* ideal lattice, i.e., an ideal lattice uniformly distributed on the Arakelov class group. In Chapter 5 the random walk theorem of the current chapter will be applied in order to ‘randomize’ a fixed ideal lattice to a uniformly random ideal lattice *without disturbing its geometrical structure too much*.

These two research directions both have an underlying question that regards

randomization of ideal lattices in the Arakelov class group, but in such a way that it does not change the nature of these ideal lattices too much. In this chapter we propose an approach to do so by means of a *random walk*, a technique that has already been deployed in various areas of mathematics. In the field of algebraic geometry, for example, one can show by a random walk technique via isogenies on elliptic curves [JMV09] (and more general abelian varieties [JW15]) that the discrete logarithm problem in a randomly chosen elliptic curve is as hard as in any other in the same isogeny class. The random walk approach on the Arakelov class group, as treated in this chapter, is heavily inspired by these results.

Related work

We note that recent works [PHS19; Lee+19] were already implicitly relying on Arakelov theory. More specifically, the lattice given in Section 3.1 of [PHS19] is precisely the lattice of slightly more general Arakelov class relations between the appropriate set of degree-zero Arakelov divisors. In fact, by extending our theorem to Arakelov divisors that include (complex) phases in the infinite places, one can obtain upper bounds for the covering radius of the relation lattices, at least for sufficiently large factor bases. With more effort one may be able to eliminate Heuristic 4 from [PHS19] or Heuristic 1 of [Lee+19].

Applications

One application of the random walk theorem concerns a worst-case to average-case reduction, as treated in Chapter 5. By uniformizing over the Arakelov class group, one can ‘randomize’ an input lattice to a uniformly random lattice, without changing the geometry of the initial lattice too much. By using the efficient machinery of random walks, one can then obtain a worst-case to average-case reduction for Hermite-SVP with only a small loss in the quality of the output.

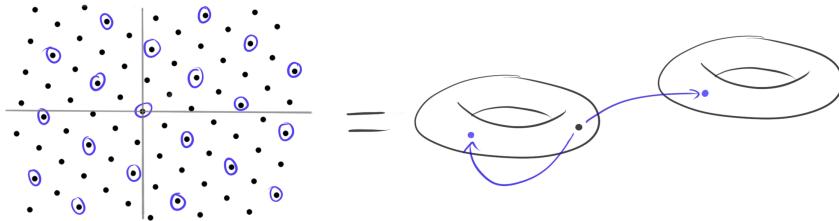
Another direct application of the random walk theorem concerns *ideal sampling* and is the object of Chapter 6. Namely, we note that many algorithms [BF14; Bia+17; BP17] rely on finding elements α in an ideal \mathfrak{a} such that $\alpha\mathfrak{a}^{-1}$ is easy to factor (e.g. prime, near-prime, or B -smooth). Such algorithms are analyzed only heuristically, by treating $\alpha\mathfrak{a}^{-1}$ as a uniformly sampled ideal, and applying known results on the density of prime or smooth ideals. The random walk theorem of this chapter allows to adjust this strategy and make the reasoning rigorous, see Chapter 6. This particular application allows to develop an efficient algorithm that computes the power residue symbol, which is the object of Chapter 7.

The result

In this chapter we show a new versatile tool: we prove that, subject to the Riemann Hypothesis for Hecke L -functions, certain random walks on the Arakelov class group have a rapid mixing property.

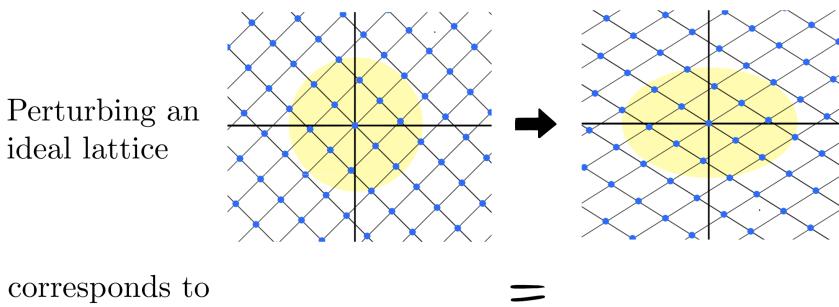
The random walk used in the result of this chapter can be seen as a combination of two different random walks, namely a discrete one and a continuous one. This is due to the fact that the Arakelov class group is topologically a disjoint union of (hyper)tori; the discrete walk ‘jumps’ from one torus to the other (see Figure 4.2), whereas the continuous walk crawls on the surface of one torus. These two different shapes of the random walk have an intuitive interpretation for ideal lattices associated with the Arakelov class group. Namely, the discrete walk corresponds to taking a random sub-ideal lattice of a given ideal lattice, also known as *sparsification*, whereas the continuous walk corresponds to *disturbing* the ideal lattice by multiplying each coordinate by a scalar.

Both the continuous and the discrete part of the random walk change the original nature of an input ideal lattice; the longer the random walk on an ideal lattice, the more disturbed this ideal lattice becomes. In the two research directions sketched earlier in this introduction, it was of fundamental importance that the final randomized ideal lattice does not *differ too much* from the input ideal lattice, but is nevertheless uniformly randomly



Taking a sub ideal lattice corresponds to a jump on the Arakelov class group.

Figure 4.2.: The discrete walk on the Arakelov class group mostly jumps from one torus to another – but it is also possible that it jumps to another distant place on the same torus.



corresponds to

a small ‘crawl’ on one single connected component of the Arakelov class group.

Figure 4.3.: The continuous walk on the Arakelov class group stays within the same torus and within a short distance of the initial point

distributed in the Arakelov class group. In other words, one would like to have an *as short as possible* random walk that still achieves uniformity in the Arakelov class group. Therefore, the study in this chapter boils down to the following succinct question.

How fast does a random walk in the Arakelov class group converge to the uniform distribution?

Concretely, the discrete walk involves so-called *finite places* and happens by multiplying the input Arakelov class by N random prime ideals the set \mathcal{P} of all prime ideals with norm bounded by B . Contrarily, the continuous walk involves the *infinite places* and happens by applying a Gaussian noise of deviation s to the input Arakelov class. Noting that the Gaussian noise of deviation s roughly covers a surface of s^r (where r is the rank of the unit group of K) and assuming that the $|\mathcal{P}|^N$ discrete jumps are sufficiently equidistributed, we can heuristically expect the random walk to yield a uniform distribution whenever $s^r \cdot |\mathcal{P}|^N \approx |\text{Pic}_K^0|$, where $|\text{Pic}_K^0|$ is the total surface of the Arakelov class group Pic_K^0 (see Figure 4.4). In fact, one can argue that this is the best situation that one can expect, due to the absence of overlapping Gaussians. This intuitive reasoning therefore can be considered as a *combined lower bound* on the parameters s , N and $|\mathcal{P}|$.

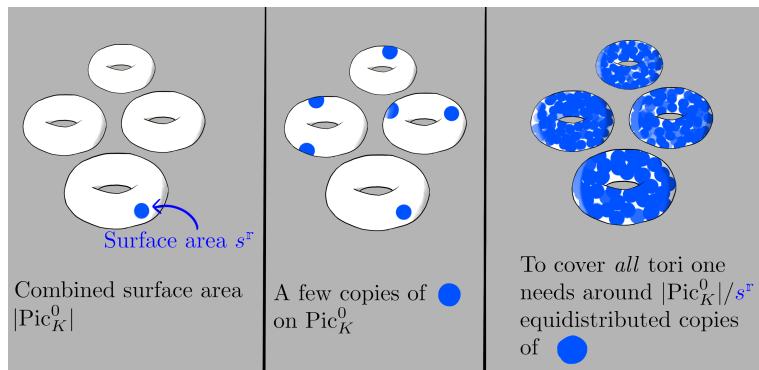


Figure 4.4.: To cover the entire Arakelov class group Pic_K^0 with Gaussians of surface area roughly s^r , we need around $|\text{Pic}_K^0|/s^r$ equidistributed copies of that Gaussian distribution.

In this work we show that the reality does not deviate much from this optimal intuitive combined lower bound on these parameters (see Theorem 4.3), assuming the Riemann Hypothesis for Hecke L-functions. The need for this specific hypothesis is not surprising: the heuristic argument assumes a reasonably controllable equidistribution of prime ideals in the Arakelov class group. For such an equidistribution of prime ideals to be within useful bounds, one often needs some form of the Riemann Hypothesis.

The actual proof that a random walk procedure yields a uniform distribution on the Arakelov class group happens by means of harmonic analysis; due to the fact that the Arakelov class group is compact and abelian, one can apply Fourier theory.

For an intuition for this proof it is convenient to consider the continuous walk before the discrete walk, i.e., we assume that we start with a reasonably narrow Gaussian distribution on one connected component of the Arakelov class group. The discrete walk, i.e., act of multiplying by a random bounded prime ideal, can be seen as a Hermitian operator on distributions on the Arakelov class group, called the *Hecke operator*. One can show that the eigenfunctions of this operator are precisely the *characters* on the Arakelov class group. Furthermore, all low-frequency characters have eigenvalues whose absolute value is sufficiently below one, except the unit character, which is kept intact under this Hecke operator. As Gaussian distributions only have a negligible contribution from high-frequency characters, applying the Hecke operator sufficiently often on this Gaussian suppresses all characters except the unit character, thus yielding an almost uniform distribution.

In this chapter we consider a slight generalization of the Arakelov class group, called the Arakelov *ray* class group with respect to a modulus ideal $\mathfrak{m} \subseteq \mathcal{O}_K$. This Arakelov ray class group is essentially an Arakelov class group where we ‘leave out’ the primes dividing the ideal \mathfrak{m} and where the principal divisors must equal 1 modulo \mathfrak{m} . This generalization is needed in Chapter 7, in which we show that the power residue symbol can be computed within zero-error probabilistic polynomial time. To recover the ordinary Arakelov class group, one just puts $\mathfrak{m} = \mathcal{O}_K$.

4.3. Random Walk Theorem for the Arakelov Ray Class Group

In this section, we prove Theorem 4.3, on random walks on the Arakelov ray class group. Starting with a point in the hyperplane $H \subseteq \text{Div}_{K^m}^0$, sampled according to a Gaussian distribution, we prove that multiplying this point sufficiently often by small random prime ideals yields a random ray divisor that is very close to uniformly distributed in the Arakelov ray class group (i.e., modulo principal ray divisors). The proof of Theorem 4.3 requires various techniques, extensively treated in Sections 4.3.2 to 4.3.7, and summarized in the following.

Hecke operators. The most important tool for proving Theorem 4.3 is that of a Hecke operator, whose definition and properties are covered in Section 4.3.2. This specific kind of operator acts on the space of probability distributions on $\text{Pic}_{K^m}^0$, and has the virtue of having the characters of $\text{Pic}_{K^m}^0$ as eigenfunctions.

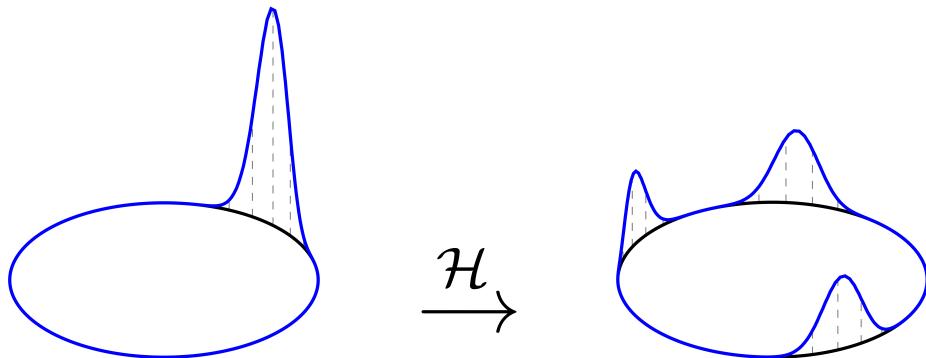


Figure 4.5.: On distributions on the Arakelov class group (here portrayed as a single circle) the Hecke operator has the effect of taking multiple shifts of the distribution and taking the average of those. In the pictured example, the Hecke operator maps the input Gaussian distribution on the circle to a distribution consisting of the average of three shifts of this Gaussian distribution.

Eigenvalues of Hecke operators. The aim of the proof is showing that applying this Hecke operator repeatedly on an appropriate initial distribution yields the uniform distribution on $\text{Pic}_{K^m}^0$. The impact of consecutive applications of the Hecke operator can be studied by considering the eigenvalues of its eigenfunctions (which are the characters of $\text{Pic}_{K^m}^0$). Classical results from analytic number theory show that the eigenvalues of these characters are (in absolute value) sufficiently smaller than 1, whenever the so-called analytic conductor of the corresponding character is not too large. An exception is the unit character, which is fixed under each Hecke operation. This classical result and how to apply it in our specific setting is covered in Section 4.3.3.

The analytic conductor. The Hecke operator thus quickly ‘damps out’ all characters with small analytic conductor (except the unit character). In Section 4.3.4, we examine which quantities of a character of $\text{Pic}_{K^m}^0$ define the analytic conductor. It turns out that this analytic conductor is closely related to how the character acts on the hypertorus T^m defined by the log ray unit lattice. The higher the frequency of this character on the hypertorus, the larger the analytic conductor. This frequency can be measured by the norm of the uniquely associated dual log ray unit lattice point of the character. In fact, we establish a bound on the analytic conductor of a character in terms of the norm of its associated dual lattice point.

Fourier analysis on the hypertorus T^m . To summarize, low-frequency (non-trivial) characters on $\text{Pic}_{K^m}^0$ (i.e., with small analytic conductor) are quickly damped out by the action of the Hecke character, whereas for high-frequency characters we do not have good guarantees on the speed at which they damp out. To resolve this issue, we choose an initial distribution whose character decomposition has only a negligible portion of high-frequency oscillatory characters. An initial distribution that nicely satisfies this condition is the Gaussian distribution (on the hypertorus). To examine the exact amplitudes of the occurring characters of this Gaussian distribution, we need Fourier analysis on this hypertorus, as covered in Section 4.3.5.

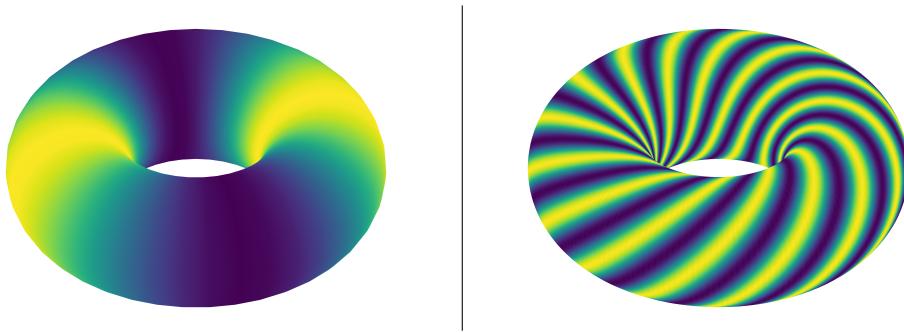


Figure 4.6.: On the left, a character with a low frequency on the hypertorus T^m is pictured. Contrarily, on the right one can see a character that has a rather high frequency on the hypertorus. As a result, the character of the left image has a *lower analytic conductor* than the character on the right image.

Splitting up the character decomposition. In this part of the proof, which is covered in Section 4.3.6, we write the Gaussian distribution into its character decomposition, where we separate the high-frequency characters, the low-frequency character and the unit character. Applying the Hecke operator often enough, damps out the low-frequency ones, and as the high-frequency characters were only negligibly present anyway, one is left with (almost only) the unit character. This corresponds to a uniform distribution.

Conclusion. By assembling all technical results and choosing appropriate parameters, we arrive at the main theorem, which is stated and proved in Section 4.3.7.

4.3.1. Main result

Definition 4.1 (Random Walk Distribution in $\text{Div}_{K^m}^0$). *For a number field K , we denote by $\mathcal{W}_{\text{Div}_{K^m}^0}(B, N, s)$ the distribution on $\text{Div}_{K^m}^0$ that is obtained by the following random walk procedure.*

Sample $x \in H \subseteq \log K_{\mathbb{R}}$ according to a centered Gaussian distribution with standard deviation s . Subsequently, sample N ideals \mathfrak{p}_j uniformly from the

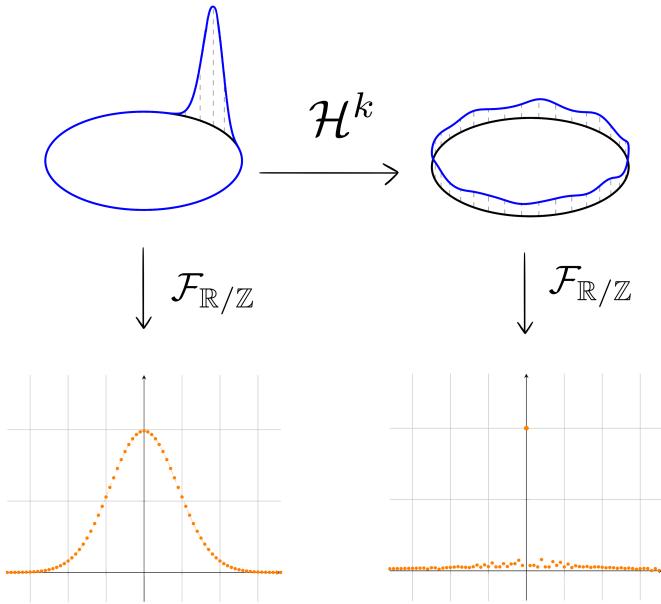


Figure 4.7.: The main theorem of this chapter is proven by resorting to Fourier analysis. The initial distribution (a Gaussian distribution in this example) on the Arakelov ray class group can be decomposed into a sum of characters. The Hecke operator \mathcal{H} has a diminishing effect on non-unit characters. Therefore, applying it sufficiently many times results in a distribution that is almost uniform.

set of all prime ideals coprime with \mathfrak{m} with norm bounded by B . Finally, output $x + \sum_{j=1}^N d^0(\mathfrak{p}_j)$, where $x \in \text{Div}_{K^{\mathfrak{m}}}^0$ is understood via the injection $H \hookrightarrow \text{Div}_{K^{\mathfrak{m}}}^0$.

Definition 4.2 (Random Walk Distribution in $\text{Pic}_{K^{\mathfrak{m}}}^0$). *By $\mathcal{W}_{\text{Pic}_{K^{\mathfrak{m}}}^0}(B, N, s)$, we denote the distribution on the Arakelov class group obtained by sampling \mathbf{a} from $\mathcal{W}_{\text{Div}_{K^{\mathfrak{m}}}^0}(B, N, s)$ and taking the Arakelov ray class $[\mathbf{a}] \in \text{Pic}_{K^{\mathfrak{m}}}^0$.*

Theorem 4.3 (Random Walks on the Arakelov Ray Class Group, ERH). *Let $\varepsilon > 0$ and $s > 0$ be any positive real numbers and let $k \in \mathbb{R}_{>0}$ be a positive real number as well. Putting $\tilde{s} = \min(\sqrt{2} \cdot s, 1/\eta_1(\Lambda_{K^{\mathfrak{m},1}}^*))$, there exists a*

4. Random Walks on Arakelov Ray Class Groups

bound $B = \tilde{O}(n^{2k}[n^2(\log \log(1/\varepsilon))^2 + n^2(\log(1/\tilde{s}))^2 + (\log(|\Delta_K| \mathcal{N}(\mathfrak{m})))^2])$ such that for any integer

$$N \geq \left\lceil \frac{1}{2k \log n} \cdot (\mathfrak{r} \cdot \log(1/\tilde{s}) + \log|\text{Pic}_{K^{\mathfrak{m}}}^0| + 2 \log(1/\varepsilon) + 2) \right\rceil, \quad (4.46)$$

the random walk distribution $\mathcal{W}_{\text{Pic}_{K^{\mathfrak{m}}}^0}(B, N, s)$ is ε -close to uniform in $L_1(\text{Pic}_{K^{\mathfrak{m}}}^0)$, i.e.,

$$\left\| \mathcal{W}_{\text{Pic}_{K^{\mathfrak{m}}}^0}(B, N, s) - \mathcal{U}(\text{Pic}_{K^{\mathfrak{m}}}^0) \right\|_1 \leq \varepsilon.$$

4.3.2. Hecke Operators

A key tool to analyze random walks on $\text{Pic}_{K^{\mathfrak{m}}}^0$ are Hecke operators, which allow to transform a given distribution into a new distribution obtained by adding one random step. This particular Hecke operator, though mainly interpreted as an operator on distributions, can in fact be applied on any function on $\text{Pic}_{K^{\mathfrak{m}}}^0$.

Definition 4.4 (The Hecke operator). *Let \mathcal{P} be a finite subset of prime ideals of the number field K not dividing the modulus \mathfrak{m} , and let $\text{Pic}_{K^{\mathfrak{m}}}^0$ be the Arakelov ray class group with respect to this modulus $\mathfrak{m} \subseteq \mathcal{O}_K$. Then we define the Hecke operator $\mathcal{H}_{\mathcal{P}} : L^2(\text{Pic}_{K^{\mathfrak{m}}}^0) \rightarrow L^2(\text{Pic}_{K^{\mathfrak{m}}}^0)$ by the following rule:*

$$\mathcal{H}_{\mathcal{P}}(f)(x) := \frac{1}{|\mathcal{P}|} \sum_{\mathfrak{p} \in \mathcal{P}} f(x - [d^0(\mathfrak{p})])$$

Concretely, the Hecke operator on the Arakelov ray class group $\text{Pic}_{K^{\mathfrak{m}}}^0$ sends a distribution over $\text{Pic}_{K^{\mathfrak{m}}}^0$ to an *average of shifts* of this distribution, see Figure 4.5.

Lemma 4.5 (Eigenfunctions of the Hecke operator). *The Hecke operator $\mathcal{H}_{\mathcal{P}} : L^2(\text{Pic}_{K^{\mathfrak{m}}}^0) \rightarrow L^2(\text{Pic}_{K^{\mathfrak{m}}}^0)$ has the characters $\chi \in \widehat{\text{Pic}_{K^{\mathfrak{m}}}^0}$ as eigenfunctions, with eigenvalues $\lambda_{\chi} = \frac{1}{|\mathcal{P}|} \sum_{\mathfrak{p} \in \mathcal{P}} \overline{\chi}([d^0(\mathfrak{p})])$, i.e.,*

$$\mathcal{H}_{\mathcal{P}}(\chi) = \lambda_{\chi} \chi.$$

Proof. Let $\chi \in \widehat{\text{Pic}}_{K^m}^0$ be a character on $\text{Pic}_{K^m}^0$. We have

$$\begin{aligned}\mathcal{H}_{\mathcal{P}}(\chi)(x) &= \frac{1}{|\mathcal{P}|} \sum_{\mathfrak{p} \in \mathcal{P}} \chi(x - [d^0(\mathfrak{p})]) \\ &= \frac{1}{|\mathcal{P}|} \sum_{\mathfrak{p} \in \mathcal{P}} \chi(x) \bar{\chi}([d^0(\mathfrak{p})]) = \left(\frac{1}{|\mathcal{P}|} \sum_{\mathfrak{p} \in \mathcal{P}} \bar{\chi}([d^0(\mathfrak{p})]) \right) \cdot \chi(x).\end{aligned}$$

So $\mathcal{H}_{\mathcal{P}}(\chi) = \lambda_{\chi} \chi$ with $\lambda_{\chi} = \frac{1}{|\mathcal{P}|} \sum_{\mathfrak{p} \in \mathcal{P}} \bar{\chi}([d^0(\mathfrak{p})])$. \square

Note that $\mathcal{H}_{\mathcal{P}}(\mathbf{1}) = \mathbf{1}$, for the trivial character $\mathbf{1} \in \widehat{\text{Pic}}_{K^m}^0$, so $\lambda_{\mathbf{1}} = 1$. For any other character χ it is evident from the above that $|\lambda_{\chi}| \leq 1$.

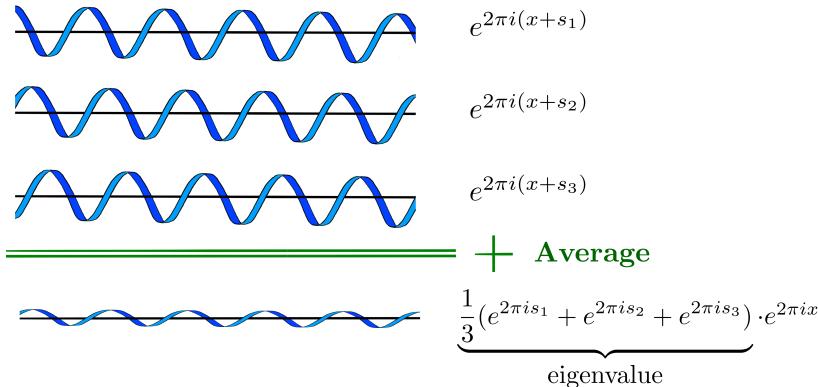


Figure 4.8.: Taking the average of shifted complex exponential functions yields a ‘flattened’ complex exponential function. Hence those complex exponentials (i.e., characters) are eigenfunctions of the Hecke operator, having eigen values in absolute value bounded by one.

4.3.3. Bounds on Eigenvalues of Hecke Operators

In the remaining part of this chapter we consider the Hecke operator whose prime set \mathcal{P} consists of *all* primes with norm bounded by B that are not dividing the modulus \mathfrak{m} . Assuming the Extended Riemann Hypothesis for

Hecke L-functions (see Definition 2.10) and using classical results from analytic number theory, one can show that the eigenvalues of these specific Hecke operators tend to zero if B grows to infinity for non-trivial characters. More specifically, omitting quantities like the conductor and the discriminant for the moment, one can show that the eigenvalues of the non-trivial characters of these Hecke operators are essentially bounded by $O(B^{-1/2})$ in a non-uniform way.

Proposition 4.6 (Bound on the eigenvalues of the Hecke operator, ERH). *Let \mathcal{P} be the set of all primes of K not dividing \mathfrak{m} and with norm bounded by $B \in \mathbb{N}$. Then, assuming the extended Riemann hypothesis (Definition 2.10), the eigenvalue λ_χ of any non-constant eigenfunction $\chi \in \widehat{\text{Pic}}_{K^{\mathfrak{m},1}}^0$ of the Hecke operator satisfies*

$$\lambda_\chi = O\left(\frac{\log(B) \log(B^n \cdot |\Delta_K| \cdot \mathcal{N}(\mathfrak{m}) \cdot \mathfrak{q}_\infty(\chi))}{B^{1/2}}\right),$$

provided that $B \geq \max((12 \log |\Delta_K| + 8n + 28)^4, 3 \cdot 10^{11}, 16 \cdot \omega(\mathfrak{m})^2)$, where $\mathfrak{q}_\infty(\chi)$ is the infinite part of the analytic conductor of the character χ , as in Definition 4.12 (cf. [IKS04, Eq. (5.6)]) and $\omega(\mathfrak{m})$ is the number of different prime ideal divisors of \mathfrak{m} .

Notation 4.7. We denote by $\mathcal{M} : \mathcal{I}_K \rightarrow \mathbb{R}_{>0}$ the von Mangoldt function for number fields K . The value $\mathcal{M}(\mathfrak{a})$ equals $\log(\mathcal{N}(\mathfrak{p}))$ whenever \mathfrak{a} is a power of a prime ideal \mathfrak{p} and zero otherwise. We also define the function $\widetilde{\mathcal{M}} : \mathcal{I}_K \rightarrow \mathbb{R}_{>0}$, for which $\widetilde{\mathcal{M}}(\mathfrak{a}) = \log(\mathcal{N}(\mathfrak{a}))$ whenever $\mathcal{N}(\mathfrak{a})$ is prime and zero otherwise.

In order to apply analytic number-theoretic results, we need to eliminate the non-split primes of the number field from the character sums arising in the eigenvalues of the Hecke operator. This happens in the following lemma, whose proof follows exactly the outline of [Wes18, Cor. 2.3.5].

Lemma 4.8. *For any character $\chi : \mathcal{I}_K^{\mathfrak{m}} \rightarrow \mathbb{C}$, we have*

$$\sum_{\substack{\mathcal{N}(\mathfrak{a}) \leq B \\ \mathfrak{a} + \mathfrak{m} = \mathcal{O}_K}} \chi(\mathfrak{a}) \mathcal{M}(\mathfrak{a}) - \sum_{\substack{\mathcal{N}(\mathfrak{a}) \leq B \\ \mathfrak{a} + \mathfrak{m} = \mathcal{O}_K}} \chi(\mathfrak{a}) \widetilde{\mathcal{M}}(\mathfrak{a}) = O(n\sqrt{B}) \quad (4.47)$$

where the sums are over all integral ideals coprime to the modulus \mathfrak{m} and with norm bounded by B .

Proof. Any nonzero entry $\chi(\mathfrak{a})[\mathcal{M}(\mathfrak{a}) - \widetilde{\mathcal{M}}(\mathfrak{a})]$ arises from an ideal \mathfrak{a} that is a power of a prime ideal and that does not have prime norm. As there are at most $n = [K : \mathbb{Q}]$ prime ideals above each prime number, we see that the left side of Equation (4.47) must be bounded by

$$\begin{aligned} n \sum_{\substack{p^\ell \leq B \\ \ell \geq 2}} \ln(p) &\leq n \sum_{\substack{p \leq \sqrt{B} \\ 2 \leq \ell \leq \frac{\ln B}{\ln p}}} \ln(p) \\ &\leq n \sum_{p \leq \sqrt{B}} \ln p \frac{\ln B}{\ln p} = n \cdot \pi(B^{1/2}) \cdot \ln B = O(n \cdot B^{1/2}), \end{aligned}$$

where π is the prime counting function over \mathbb{Z} and where the last bound is obtained by the prime number theorem (see Theorem 2.12). \square

Proof of Proposition 4.6. Assuming the Extended Riemann Hypothesis, we have the following classical analytic result¹ [IKS04, Thm. 5.15] for any non-trivial character $\chi \in \widehat{\text{Pic}}_{K^{\mathfrak{m}}}^0$.

¹Any character on the Arakelov ray class group can be seen as a Hecke character, by projecting the idèle class group to the Arakelov ray class group. Since characters χ on the Arakelov ray class group are defined on $\mathcal{I}_K^{\mathfrak{m}}$, the conductor \mathfrak{f}_χ divides \mathfrak{m} . The analytic conductor $\mathfrak{q}(\chi)$ is then equal to $|\Delta_K| \cdot \mathcal{N}(\mathfrak{f}_\chi) \cdot \mathfrak{q}_\infty(\chi) \leq |\Delta_K| \cdot \mathcal{N}(\mathfrak{m}) \cdot \mathfrak{q}_\infty(\chi)$, where Δ_K is the discriminant of the number field K and $\mathfrak{q}_\infty(\chi)$ is the infinite part of the analytic conductor; see, for example, [IKS04, p. 129 & Eq. (5.7)].

The phrasing of the theorem in Iwaniec & Kowalski [IKS04, Thm. 5.15] involves the function $\psi(f, x) = \sum_{n \leq x} \Lambda_f(n)$, defined in [IKS04, Eq. (5.46), p. 110], where $\Lambda_f(n)$ is supported only on prime powers and arises from $-L'(f, s)/L(f, s) = \sum_{n \geq 1} \Lambda_f(n) n^{-s}$ [IKS04, Eq. (5.25), p. 102]. In our case, $f = \chi$ is a Hecke character with conductor \mathfrak{m} , which means that the associated L -function avoids \mathfrak{m} : $L(s, \chi) = \prod_{\mathfrak{p} \mid \mathfrak{f}_\chi} (1 - \chi(\mathfrak{p}) \mathcal{N}(\mathfrak{p})^{-s})^{-1}$ (see [Lan12, Ch. XIV, §8, p. 299]). By taking the logarithmic derivative of $L(\chi, s)$ one obtains that $\Lambda_\chi(n) = \sum_{\substack{\mathcal{N}(\mathfrak{a})=n \\ \mathfrak{a}+\mathfrak{m}=\mathcal{O}_K}} \chi(\mathfrak{a}) \mathcal{M}(\mathfrak{a})$, where \mathcal{M} is the van Mangoldt function as in Notation 4.7. The same theorem [IKS04, Thm. 5.15] involves a number r indicating the order of the pole of zero at $s = 1$ of the respective L -function. In the case of non-trivial Hecke characters χ this order r is zero, see [IKS04, Ch. 5, p. 94 and p. 129] or [Lan12, Ch. XV, §4, Thm. 2].

$$\sum_{\substack{\mathcal{N}(\mathfrak{a}) \leq B \\ \mathfrak{a} + \mathfrak{m} = \mathcal{O}_K}} \mathcal{M}(\mathfrak{a}) \chi(\mathfrak{a}) = O(B^{1/2} \log(B) \log(B^n \cdot |\Delta_K| \cdot \mathcal{N}(\mathfrak{m}) \cdot \mathfrak{q}_\infty(\chi))),$$

where $\mathfrak{q}_\infty(\chi)$ is the infinite part of the analytic conductor of χ , and where \mathcal{M} is the von Mangoldt function for the number field K (see Notation 4.7). In this expression, and in all subsequent expressions in this proof, the summation is over integral ideals \mathfrak{a} *coprime* with the modulus \mathfrak{m} , as indicated by the phrase ‘ $\mathfrak{a} + \mathfrak{m} = \mathcal{O}_K$ ’.

According to Lemma 4.8, the sums

$$\sum_{\substack{\mathcal{N}(\mathfrak{a}) \leq B \\ \mathfrak{a} + \mathfrak{m} = \mathcal{O}_K}} \chi(\mathfrak{a}) \widetilde{\mathcal{M}}(\mathfrak{a}) \quad \text{and} \quad \sum_{\substack{\mathcal{N}(\mathfrak{a}) \leq B \\ \mathfrak{a} + \mathfrak{m} = \mathcal{O}_K}} \chi(\mathfrak{a}) \mathcal{M}(\mathfrak{a})$$

(over integral ideals coprime with \mathfrak{m}) differ by at most $O(nB^{1/2})$, and therefore

$$\begin{aligned} A(B) &:= \sum_{2 \leq j \leq B} a_j = \sum_{\substack{\mathcal{N}(\mathfrak{a}) \leq B \\ \mathfrak{a} + \mathfrak{m} = \mathcal{O}_K}} \chi(\mathfrak{a}) \widetilde{\mathcal{M}}(\mathfrak{a}) \\ &= O(B^{1/2} \log(B) \log(B^n \cdot |\Delta_K| \cdot \mathcal{N}(\mathfrak{m}) \cdot \mathfrak{q}_\infty(\chi))), \end{aligned}$$

where $a_n = \sum_{\substack{\mathcal{N}(\mathfrak{a})=n \\ \mathfrak{a} + \mathfrak{m} = \mathcal{O}_K}} \chi(\mathfrak{a}) \widetilde{\mathcal{M}}(n)$ and where $\widetilde{\mathcal{M}}(n) = \log n$ whenever n is prime and zero otherwise. Using the Abel partial summation formula, and temporarily denoting $C = |\Delta_K| \cdot \mathcal{N}(\mathfrak{m}) \cdot \mathfrak{q}_\infty(\chi)$ for the sake of brevity, we deduce

$$\begin{aligned} \sum_{\substack{\mathcal{N}(\mathfrak{a}) \leq B \\ \mathfrak{a} + \mathfrak{m} = \mathcal{O}_K}} \chi(\mathfrak{p}) &= \sum_{n \leq B} a_n \frac{1}{\log n} = \frac{A(B)}{\log B} + \int_2^B A(t) \frac{dt}{t \log^2(t)} \\ &= O(B^{1/2} \log(B^n \cdot C)) + O\left(\int_2^B \frac{\log(t^n \cdot C)}{\log(t) t^{1/2}} dt\right) \\ &= O(B^{1/2} \log(B^n \cdot \underbrace{|\Delta_K| \cdot \mathcal{N}(\mathfrak{m}) \cdot \mathfrak{q}_\infty(\chi)}_C)) \end{aligned} \tag{4.48}$$

where the last equality uses the fact that

$$\int_2^B \frac{\log(t^n \cdot C)}{\log(t)t^{1/2}} dt \leq \log(B^n \cdot C) / \log(2) \int_2^B t^{-1/2} dt = O(B^{1/2} \cdot \log(B^n \cdot C)).$$

As the composition $\chi \circ [d^0(\cdot)] : \mathcal{I}_K^{\mathfrak{m}} \rightarrow \mathbb{C}$ is a Hecke character on ideals coprime to \mathfrak{m} , and $|\mathcal{P}| = \Theta(B/\log(B))$ (see² Lemma 2.13), we apply Equation (4.48) to obtain

$$\begin{aligned} \lambda_{\chi} &= \frac{1}{|\mathcal{P}|} \sum_{\mathfrak{p} \in \mathcal{P}} \bar{\chi}(d^0(\mathfrak{p})) = \frac{1}{|\mathcal{P}|} O(B^{1/2} \log(B^n \cdot |\Delta_K| \cdot \mathcal{N}(\mathfrak{m}) \cdot \mathfrak{q}_{\infty}(\chi))) \\ &= O\left(B^{-1/2} \log(B) \log(B^n \cdot |\Delta_K| \cdot \mathcal{N}(\mathfrak{m}) \cdot \mathfrak{q}_{\infty}(\chi))\right) \end{aligned}$$

which finishes the proof. \square

4.3.4. The Infinite Analytic Conductor

In the bounds of Section 4.3.3, the infinite analytic conductor $\mathfrak{q}_{\infty}(\chi)$ of a character $\chi : \text{Pic}_{K^{\mathfrak{m}}}^0 \rightarrow \mathbb{C}$ plays a large role. In this section, we show that this infinite analytic conductor $\mathfrak{q}_{\infty}(\chi)$ is closely related to the dual logarithmic ray unit lattice point $\ell^* \in \Lambda_{K^{\mathfrak{m}}, 1}^* = \text{Log}(\mathcal{O}_{K^{\mathfrak{m}}, 1}^*)^*$ that is uniquely associated with the character $\chi|_{T^{\mathfrak{m}}} : T^{\mathfrak{m}} \rightarrow \mathbb{C}$. We analyze the infinite analytic conductor of a character $\chi : \text{Pic}_{K^{\mathfrak{m}}}^0 \rightarrow \mathbb{C}$ using the following facts:

- Any Arakelov ray class group $\text{Pic}_{K^{\mathfrak{m}}}^0$ is a quotient group of the degree-zero idèle class group \mathcal{C}_K^0 . We will prove that there is a canonical projection $\mathcal{C}_K^0 \rightarrow \text{Pic}_{K^{\mathfrak{m}}}^0$ for all integral moduli $\mathfrak{m} \in \mathcal{I}_K$. This immediately has as a consequence that any character $\text{Pic}_{K^{\mathfrak{m}}}^0 \rightarrow \mathbb{C}$ yields an induced character on the idèle class group \mathcal{C}_K^0 by precomposition with this projection. Summarizing: Any character on $\text{Pic}_{K^{\mathfrak{m}}}^0$ is a Hecke character.
- There is a canonical map $K_{\mathbb{R}}^0 \rightarrow \mathcal{C}_K^0$, so any Hecke character $\chi : \mathcal{C}_K^0 \rightarrow \mathbb{C}$ induces a derived character $K_{\mathbb{R}}^0 \rightarrow \mathcal{C}_K^0 \xrightarrow{\chi} \mathbb{C}$. Characters on the group $K_{\mathbb{R}}^0$ are known to have a very specific shape, which can be described

²For this to be true, the lower bound on $B \geq \max((12 \log |\Delta_K| + 8n + 28)^4, 3 \cdot 10^{11}, 16 \cdot \omega(\mathfrak{m})^2)$ is needed, see Lemma 2.13.

in terms of an $2(n_{\mathbb{C}} + n_{\mathbb{R}})$ -dimensional real vector. The entries of this specific vector are called the *local parameters at infinity* of the Hecke character χ .

- *Hecke characters derived from a character on the Arakelov ray class group $\text{Pic}_{K^{\mathfrak{m}}}^0$ have local parameters at infinity that are closely related to the dual lattice of the logarithmic ray unit lattice $\Lambda_{K^{\mathfrak{m},1}} = \text{Log}(\mathcal{O}_{K^{\mathfrak{m},1}}^{\times}) = \text{Log}(\mathcal{O}_K^{\times} \cap K^{\mathfrak{m},1})$. It turns out that a character on the Arakelov ray class group $\text{Pic}_{K^{\mathfrak{m}}}^0$ induces a character on $K_{\mathbb{R}}^0$ that does not depend on the *phases* of the complex numbers involved, i.e., it solely depends on the absolute values of the entries in $K_{\mathbb{R}}^0$. This means that the induced character on $K_{\mathbb{R}}^0$ factors through the logarithmic image of the absolute values of $K_{\mathbb{R}}^0/\mathcal{O}_K^{\times}$, which equals $H/\Lambda_{K^{\mathfrak{m},1}} = T^{\mathfrak{m}}$. Characters on this *ray unit torus* $T^{\mathfrak{m}}$ are uniquely described by a dual logarithmic unit lattice point $\ell^* \in \Lambda_{K^{\mathfrak{m},1}}^*$.*
- *The infinite analytic conductor \mathfrak{q}_{∞} is just a specific product of the local parameters of χ over all infinite places ν . More specifically, we have*

$$\mathfrak{q}_{\infty}(\chi) = \prod_{\nu \text{ real}} (3 + ||n_{\nu}| + i\phi_{\nu}|) \cdot \prod_{\nu \text{ complex}} (3 + |n_{\nu} + i\phi_{\nu}|)(3 + ||n_{\nu}| + i\phi_{\nu} + 1|).$$

Via this formula, one can relate the size of the infinite analytic conductor $\mathfrak{q}_{\infty}(\chi)$ with the length of the dual logarithmic ray unit lattice.

In the following text we will elaborate on these four facts, for each fact a paragraph.

The Arakelov ray class group $\text{Pic}_{K^{\mathfrak{m}}}^0$ is a quotient group of \mathcal{C}_K^0

Let $\mathcal{J}_{K^{\mathfrak{m}}} \subseteq \mathcal{J}_K$ be the subgroup of idèles that satisfy $a_{\nu} \equiv 1$ modulo $\mathfrak{p}_{\nu}^{\text{ord}_{\nu}(\mathfrak{m})}$ for any place ν with $\mathfrak{p}_{\nu} \mid \mathfrak{m}$. More precisely,

$$\mathcal{J}_{K^{\mathfrak{m}}} = \{(a_{\nu})_{\nu} \in \mathcal{J}_K \mid a_{\nu} \in 1 + \mathfrak{p}_{\nu}^{\text{ord}_{\nu}(\mathfrak{m})} \text{ for all places } \nu \text{ with } \mathfrak{p}_{\nu} \mid \mathfrak{m}\}$$

Via the inclusion $\mathcal{J}_{K^{\mathfrak{m}}} \subseteq \mathcal{J}_K$, we have the isomorphism $\mathcal{J}_{K^{\mathfrak{m}}}/K^{\mathfrak{m},1} \xrightarrow{\sim} \mathcal{J}_K/K^*$ [Lan12, Ch. VII, §3]. The following map is surjective and has $K^{\mathfrak{m},1}$

in its kernel, which proves that the group Pic_{K^m} is a quotient group of $\mathcal{J}_{K^m}/K^{m,1}$ and therefore of \mathcal{J}_K/K^* as well.

$$\mathcal{J}_{K^m} \rightarrow \text{Pic}_{K^m}, (a_\nu)_\nu \longmapsto \sum_{\nu \nmid \infty} \text{ord}_{\mathfrak{p}_\nu}(a_\nu) \cdot (\mathfrak{p}_\nu) + \sum_{\nu \mid \infty} [K_\nu : \mathbb{R}] \cdot \log |a_\nu| \cdot (\nu).$$

Here, we mean by $\nu \mid \infty$ that ν is an infinite place (i.e., associated with an embedding $K \hookrightarrow \mathbb{C}$) and by $\nu \nmid \infty$ that ν is a finite place (i.e., associated with a prime ideal \mathfrak{p}). The degree maps $\deg : \mathcal{J}_K \rightarrow \mathbb{R}_{>0}, (a_\nu)_\nu \mapsto \prod_\nu |a_\nu|_\nu$ and $\deg : \text{Pic}_{K^m} \rightarrow \mathbb{R}_{>0}, \sum_{\mathfrak{p} \nmid \mathfrak{m}} n_{\mathfrak{p}}(\mathfrak{p}) + \sum_{\nu \mid \infty} x_\nu(\nu) \mapsto \prod_{\mathfrak{p} \nmid \mathfrak{m}} \mathcal{N}(\mathfrak{p})^{n_{\mathfrak{p}}} \prod_\nu e^{-x_\nu}$ are compatible with each other, which implies an induced surjective map $\mathcal{C}_K^0 \rightarrow \text{Pic}_{K^m}^0$, proving that the Arakelov class group is a quotient of the degree zero idèle class group.

Any Hecke character $\chi : \mathcal{C}_K^0 \rightarrow \mathbb{C}$ induces a derived character on $K_{\mathbb{R}}^0$.

Let $\chi : \mathcal{C}_K^0 \rightarrow \mathbb{C}$ be a Hecke character. Recall that $K_{\mathbb{R}}^0 \simeq \{(x_\nu)_\nu \in K_{\mathbb{R}} \mid \prod_\nu |x_\nu|_\nu = 1\}$ (see Equation (2.10)), which embeds canonically into \mathcal{J}_K^0 ; we have the injection

$$K_{\mathbb{R}}^0 \hookrightarrow \mathcal{J}_K^0, \quad (x_\sigma)_\sigma \longmapsto (a_\nu)_{\nu_\sigma} \quad \text{where} \quad a_\nu = \begin{cases} x_\nu & \text{for } \nu \mid \infty \\ 1 & \text{for } \nu \nmid \infty \end{cases}.$$

So any character $\chi : \mathcal{C}_K^0 \rightarrow \mathbb{C}$ induces a character on $K_{\mathbb{R}}^0$ by precomposition with $K_{\mathbb{R}}^0 \rightarrow \mathcal{J}_K^0 \rightarrow \mathcal{C}_K^0 = \mathcal{J}_K^0/K^*$, which will be denoted by $\chi|_{K_{\mathbb{R}}^0}$. It is a well-known fact that any character on $K_{\mathbb{R}}^0$ is of the following shape, and is uniquely determined in that way (see [NS13, Ch. XII, Prop. 6.7]):

$$\chi : K_{\mathbb{R}}^0 \rightarrow S^1, (x_\nu)_\nu \longmapsto \prod_{\nu \mid \infty} \left(\frac{x_\nu}{|x_\nu|} \right)^{n_\nu} e^{i \cdot [K_\nu : \mathbb{R}] \cdot \phi_\nu \cdot \log |x_\nu|}, \quad (4.49)$$

with $n_\nu \in \mathbb{Z}$ if ν is complex, and $n_\nu \in \{0, 1\}$ if ν is real and $\phi_\nu \in \mathbb{R}$. Note that $[K_\nu : \mathbb{R}] = 1$ if ν is real and $[K_\nu : \mathbb{R}] = 2$ if ν is a complex embedding.

Definition 4.9. Let $\chi : \mathcal{C}_K^0 \rightarrow \mathbb{C}$ be a Hecke character and let $\chi|_{K_{\mathbb{R}}^0} : K_{\mathbb{R}}^0 \rightarrow \mathbb{C}$ be the induced character by precomposing with the map $K_{\mathbb{R}}^0 \rightarrow \mathcal{C}_K^0$. Then

$(n_\nu)_\nu \in \mathbb{Z}^{n_{\mathbb{C}}} \times \{0, 1\}^{n_{\mathbb{R}}}$ and $(\phi_\nu)_\nu \in \mathbb{R}^{n_{\mathbb{C}} + n_{\mathbb{R}}}$ for $\nu \mid \infty$ that occur when writing $\chi|_{K_{\mathbb{R}}^0}$ in the shape of Equation (4.49), are called the local parameters at infinity of χ .

The local parameters at infinity of a Hecke character on the Arakelov ray class group $\text{Pic}_{K^{\mathfrak{m}}}^0$ can be seen as a scaled point on the dual logarithmic ray unit lattice $\Lambda_{K^{\mathfrak{m}}, 1}^*$

For any character $\chi : \text{Pic}_{K^{\mathfrak{m}}}^0 \rightarrow \mathbb{C}$, we thus have an derived character $\chi|_{K_{\mathbb{R}}^0}$ on $K_{\mathbb{R}}^0$ as follows.

$$K_{\mathbb{R}}^0 \rightarrow \mathcal{C}_K^0 \rightarrow \text{Pic}_{K^{\mathfrak{m}}}^0 \xrightarrow{\chi} \mathbb{C}.$$

Of this derived character $\chi|_{K_{\mathbb{R}}^0}$ we would like to study the local parameters as in Equation (4.49). The combined map $K_{\mathbb{R}}^0 \rightarrow \mathcal{C}_K^0 \rightarrow \text{Pic}_{K^{\mathfrak{m}}}^0$ can be described by the following rule:

$$K_{\mathbb{R}}^0 \rightarrow \text{Pic}_{K^{\mathfrak{m}}}^0, (x_\nu)_\nu \longmapsto \sum_{\nu \mid \infty} [K_\nu : \mathbb{R}] \cdot \log |x_\nu| \cdot (\nu) \mod K^{\mathfrak{m}, 1}$$

So the derived character $\chi|_{K_{\mathbb{R}}^0} : K_{\mathbb{R}}^0 \rightarrow \text{Pic}_{K^{\mathfrak{m}}}^0 \rightarrow \mathbb{C}$ cannot depend on the phases of $(x_\nu)_\nu \in K_{\mathbb{R}}^0$, which means that $n_\nu = 0$ in Equation (4.49) for all $\nu \mid \infty$ for such $\chi \in \text{Pic}_{K^{\mathfrak{m}}}^0$. Also, by the fact that $\text{Pic}_{K^{\mathfrak{m}}}^0 = \text{Div}_{K^{\mathfrak{m}}}^0 / K^{\mathfrak{m}, 1}$, we have that $\mathcal{O}_K^\times \cap K^{\mathfrak{m}, 1} = \mathcal{O}_{K^{\mathfrak{m}}, 1}^\times \subseteq K_{\mathbb{R}}^0$ must lie in the kernel of $\chi|_{K_{\mathbb{R}}^0}$.

Summarizing, for characters $\chi : \text{Pic}_{K^{\mathfrak{m}}}^0 \rightarrow \mathbb{C}$, we have

$$\chi|_{K_{\mathbb{R}}^0}((x_\nu)_\nu) = \prod_{\nu \mid \infty} e^{i \cdot [K_\nu : \mathbb{R}] \cdot \phi_\nu \cdot \log |x_\nu|} = \exp \left(i \sum_{\sigma : K \rightarrow \mathbb{C}} \phi_{\nu_\sigma} \cdot \log |\sigma(x)| \right)$$

and $\chi|_{K_{\mathbb{R}}^0}(\mathcal{O}_{K^{\mathfrak{m}}, 1}^\times) = 1$. In above expression, the sum is over all embeddings $\sigma : K \rightarrow \mathbb{C}$, and ν_σ is the place ν uniquely associated with the embedding σ . This means that the following inner product satisfies

$$\left\langle \frac{1}{2\pi} (\phi_{\nu_\sigma})_\sigma, (\log |\sigma(\eta)|)_\sigma \right\rangle = \frac{1}{2\pi} \sum_{\sigma : K \rightarrow \mathbb{C}} \phi_{\nu_\sigma} \cdot \log |\sigma(\eta)| \in \mathbb{Z} \text{ for all } \eta \in \mathcal{O}_{K^{\mathfrak{m}}, 1}^\times.$$

Recalling that $\Lambda_{K^{\mathfrak{m}}, 1} = \text{Log}(\mathcal{O}_{K^{\mathfrak{m}}, 1}^\times) = \text{Log}(\mathcal{O}_K^\times \cap K^{\mathfrak{m}, 1})$, this is equivalent to $(\phi_{\nu_\sigma})_\sigma \in 2\pi\Lambda_{K^{\mathfrak{m}}, 1}^*$; i.e., the local parameters $(\phi_{\nu_\sigma})_\sigma$ are equal to 2π times

a dual logarithmic ray unit lattice point in $\Lambda_{K^m,1}^*$. Thus we proved the following lemma.

Lemma 4.10 (Local parameters of a Hecke character on the Arakelov ray class group $\text{Pic}_{K^m}^0$). *Let $\chi : \text{Pic}_{K^m}^0 \rightarrow \mathbb{C}$ be a character on the Arakelov ray class group. Then the local parameters at infinity of χ as in Definition 4.9, satisfy*

- $n_\nu = 0$ for all $\nu \mid \infty$.
- There exists a dual logarithmic ray unit lattice point $\ell^* \in \Lambda_{K^m,1}^*$ such that $\phi_\nu = \ell_{\sigma_\nu}^*$.

Corollary 4.11. *Let $\chi : \text{Pic}_{K^m}^0 \rightarrow \mathbb{C}$ be a character on the Arakelov ray class group. Then, there exists a $\ell^* \in \Lambda_{K^m,1}^* \subseteq H$ such that*

$$\chi|_{K_{\mathbb{R}}^0}((x_\nu)_\nu) = \exp(2\pi i \cdot \langle \ell^*, (x_{\nu_\sigma})_\sigma \rangle).$$

The converse is also true. For every $\ell^ \in \Lambda_{K^m,1}^*$ there exists a character $\chi : \text{Pic}_{K^m}^0 \rightarrow \mathbb{C}$ such that*

$$\chi|_{K_{\mathbb{R}}^0}((x_\nu)_\nu) = \exp(2\pi i \cdot \langle \ell^*, (x_{\nu_\sigma})_\sigma \rangle). \quad (4.50)$$

Proof. The first claim directly follows from Lemma 4.10. The second claim can be verified by the fact that one can construct a character $\chi : \text{Pic}_{K^m}^0 \rightarrow T^m \rightarrow \mathbb{C}$ that factors through the ray unit torus by the canonical quotient map $\text{Pic}_{K^m}^0 \rightarrow T^m$ of Figure 2.8. On this ray unit torus T^m it has the values induced by Equation (4.50). In fact, by multiplying this character by other characters $\chi' \in \text{Pic}_{K^m}^0$ for which $\chi'|_{T^m}$ is trivial, one obtains all characters on $\text{Pic}_{K^m}^0$ satisfying Equation (4.50). \square

The infinite analytic conductor is a product of local parameters

Definition 4.12 (Infinite analytic conductor of a Hecke character). *Let $\chi \in \widehat{\text{Pic}_{K^m}^0}$ be a character with local parameters at infinity n_ν and ϕ_ν as in*

Definition 4.9, where ν ranges over the infinite places of K . Then, we define the infinite part of the analytic conductor to be

$$\mathfrak{q}_\infty(\chi) = \prod_{\nu \text{ real}} (3 + |n_\nu| + i\phi_\nu|) \prod_{\nu \text{ complex}} (3 + |n_\nu + i\phi_\nu|)(3 + |n_\nu + i\phi_\nu + 1|) \quad (4.51)$$

Remark 4.13. The above definition of the infinite analytic conductor is obtained from [IKS04, p. 95, Eq. (5.6) with $s = 0$], where it is described in a slightly different form. In [IKS04], the functional equation lacks the complex L -functions $L_{\mathbb{C}}$. Instead, those are replaced by $L_{\mathbb{R}}(s)L_{\mathbb{R}}(s+1) = L_{\mathbb{C}}(s)$ (see [NS13, Ch. 7, Prop. 4.3(iv)]). This means that the local parameters $\kappa_\sigma, \kappa_{\bar{\sigma}}$ as in [IKS04, p. 93, Eq. (5.3)] must equal $k_\nu, k_\nu + 1$ for the embeddings $\{\sigma, \bar{\sigma}\}$ associated with the complex place ν (cf. [IKS04, p. 125]).

Lemma 4.14. Let $\mathfrak{q}_\infty(\chi)$ be the infinite part of the analytic conductor of the character $\chi \in \widehat{\text{Pic}}_{K^m}^0$, and let $\ell^* \in \Lambda_{K^m, 1}^*$ be such that $\chi|_{T^m} = \chi_{\ell^*}$, where $\Lambda_{K^m, 1}^*$ is the dual lattice of the logarithmic ray unit lattice. Then we have

$$\mathfrak{q}_\infty(\chi) \leq (4 + 2\pi \|\ell^*\| / \sqrt{n})^n$$

Proof. Let $|\ell^*|$ denote the vector ℓ^* where all entries are replaced by their absolute value. Then, by applying subsequently the triangle inequality, the norm inequality between $\|\cdot\|_1$ and $\|\cdot\|_2$ and the arithmetic-geometric mean inequality, one obtains

$$\begin{aligned} 4\sqrt{n} + 2\pi \|\ell^*\|_2 &\geq \|4 + 2\pi|\ell^*|\|_2 \geq \frac{1}{\sqrt{n}} \|4 + 2\pi|\ell^*|\|_1 \\ &\geq \sqrt{n} \left(\prod_{\sigma} (4 + 2\pi|\ell_{\sigma}^*|) \right)^{1/n} \geq \sqrt{n} \cdot \mathfrak{q}_\infty(\chi)^{1/n}. \end{aligned}$$

Dividing by \sqrt{n} and raising to the power n yields the claim. The last inequality follows just from Equation (4.51), in which $n_\nu = 0$ for all infinite ν . \square

4.3.5. Fourier Analysis on the Ray Unit Torus

Definition 4.15. Let $H \subseteq \text{Log } K_{\mathbb{R}}$ be the ambient vector space of the log ray unit lattice $\Lambda_{K^{\mathfrak{m}},1} = \text{Log}(\mathcal{O}_{K^{\mathfrak{m}},1}^{\times})$, where $\mathcal{O}_{K^{\mathfrak{m}},1}^{\times} = \mathcal{O}_K^{\times} \cap K^{\mathfrak{m},1}$. Recall the Gaussian function $\rho_s : H \rightarrow \mathbb{R}$, $x \mapsto e^{-\pi\|x\|^2/s^2}$. Denoting $T^{\mathfrak{m}} = H/\Lambda_{K^{\mathfrak{m}},1}$, we put $\rho_s|^{T^{\mathfrak{m}}} : T^{\mathfrak{m}} \rightarrow \mathbb{R}$, $x \mapsto \sum_{\ell \in \Lambda_{K^{\mathfrak{m}},1}} \rho_s(x + \ell)$.

As we have (see Lemma A.3) $\|s^{-\mathfrak{r}} \rho_s\|_{H,1} = \int_H s^{-\mathfrak{r}} \rho_s(x) dx = 1$, and

$$\left\| s^{-\mathfrak{r}} \rho_s|^{T^{\mathfrak{m}}} \right\|_{T^{\mathfrak{m}},1} = \int_{T^{\mathfrak{m}}} s^{-\mathfrak{r}} \rho_s|^{T^{\mathfrak{m}}}(x) dx = 1,$$

both functions $s^{-\mathfrak{r}} \rho_s$ and $s^{-\mathfrak{r}} \rho_s|^{T^{\mathfrak{m}}}$ can be seen as probability distributions on their respective domains \mathbb{R}^m and $T^{\mathfrak{m}}$.

Lemma 4.16 (Fourier coefficients of the periodized Gaussian). *The periodized Gaussian function $s^{-\mathfrak{r}} \rho_s|^{T^{\mathfrak{m}}} \in L^2(T^{\mathfrak{m}})$ satisfies*

$$s^{-\mathfrak{r}} \rho_s|^{T^{\mathfrak{m}}} = \sum_{\ell^* \in \Lambda_{K^{\mathfrak{m}},1}^*} a_{\ell^*} \chi_{\ell^*} \quad (4.52)$$

where $a_{\ell^*} = \frac{1}{\text{Vol}(T^{\mathfrak{m}})} \rho_{1/s}(\ell^*)$, where $\Lambda_{K^{\mathfrak{m}},1}^*$ is the dual lattice of the log unit lattice $\Lambda_{K^{\mathfrak{m}},1}$, and where $\chi_{\ell^*}(x) = e^{-2\pi i \langle x, \ell^* \rangle}$.

Proof. We have $\langle \chi_{\ell_1^*}, \chi_{\ell_2^*} \rangle = \text{Vol}(T^{\mathfrak{m}}) \cdot \delta_{\ell_1^*, \ell_2^*}$, where δ is the Kronecker delta function, i.e., $\delta_{\ell_1^*, \ell_2^*}$ equals one if $\ell_1^* = \ell_2^*$ and zero otherwise. Identifying $\widehat{T^{\mathfrak{m}}}$ and $\Lambda_{K^{\mathfrak{m}},1}^*$ via the map $\chi_{\ell^*} \mapsto \ell^*$, taking a fundamental domain F of $\Lambda_{K^{\mathfrak{m}},1}$ and spelling out the definition of $\rho_s|^{T^{\mathfrak{m}}}$, we obtain, for all $\ell^* \in \Lambda_{K^{\mathfrak{m}},1}^*$,

$$\begin{aligned} a_{\ell^*} &= \frac{1}{\text{Vol}(T^{\mathfrak{m}})} \langle s^{-\mathfrak{r}} \rho_s|^{T^{\mathfrak{m}}}, \chi_{\ell^*} \rangle \\ &= \frac{1}{\text{Vol}(T^{\mathfrak{m}})} \int_{x \in F} \sum_{\ell \in \Lambda_{K^{\mathfrak{m}},1}} s^{-\mathfrak{r}} \rho_s(x + \ell) \overline{\chi_{\ell^*}(x)} dx \\ &= \frac{1}{\text{Vol}(T^{\mathfrak{m}})} \int_{x \in H} s^{-\mathfrak{r}} \rho_s(x) \overline{\chi_{\ell^*}(x)} dx = \frac{1}{\text{Vol}(T^{\mathfrak{m}})} \mathcal{F}_H\{s^{-\mathfrak{r}} \rho_s\}(\ell^*) \\ &= \frac{1}{\text{Vol}(T^{\mathfrak{m}})} \rho_{1/s}(\ell^*). \end{aligned}$$

The last equality can be derived from the properties of the Gaussian function in Lemma 2.23. \square

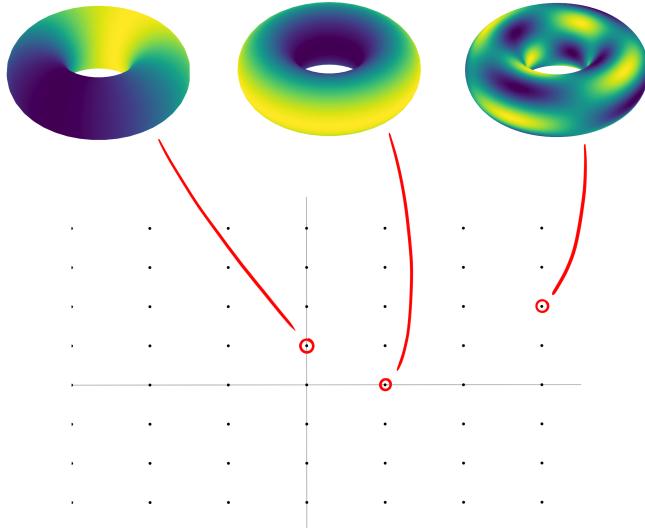


Figure 4.9.: In this picture three examples of characters ('eigenfunctions') on the ray unit torus are portrayed, together with their associated lattice points in the dual ray unit lattice at the bottom. In this particular example, one can see that characters that only depend on one rotational axis (the left two examples) have an associated dual lattice point on the x -axis or the y -axis. The 'mixed' character is associated with a dual lattice point that has both a non-zero x and y component.

4.3.6. Splitting up the Character Decomposition

Decomposing into characters on $\text{Pic}_{K^m}^0$

In Equation (4.52), the distribution $s^{-r} \rho_s|^{T^m}$ is decomposed into characters on the unit torus T^m . In order to apply the analytic bound on the eigenvalues of the Hecke operator as in Proposition 4.6, we need to decompose this distribution into characters on the *Arakelov ray class group* $\text{Pic}_{K^m}^0$ instead.

To do so, we use the identity

$$\chi_{\ell^*} = \frac{1}{|\text{Cl}_K^m|} \sum_{\substack{\chi' \in \widehat{\text{Pic}}_K^0 \\ \chi'|_{T^m} = \chi_{\ell^*}}} \chi', \quad (4.53)$$

which holds for any $\ell^* \in \Lambda_{K^m,1}^*$, where $|\text{Cl}_K^m|$ is the cardinality of the ray class group. Equation (4.53) is an identity of functions on the Arakelov ray class group Pic_K^0 , where χ_{ℓ^*} is defined to be zero everywhere, except on the torus $T^m \subseteq \text{Pic}_K^0$, the original domain of the function χ_{ℓ^*} . In this identity, χ' ranges over all characters $\chi' \in \widehat{\text{Pic}}_K^0$ which are identical to χ_{ℓ^*} when restricted to the unit group torus T^m . These characters χ' are called the *extensions* of the character χ_{ℓ^*} with respect to Pic_K^0 , and it can be shown that for each $\ell^* \in \Lambda_{K^m,1}^*$ there are exactly $|\text{Cl}_K^m|$ such extensions (see [DE16, Cor. 3.6.2]).

The identity in Equation (4.53) follows essentially from the same argument that is used to prove general character orthogonality properties (see [Ser77, §2.3]).

Splitting up the character decomposition in a low-frequency and a high-frequency part

By above reasoning, we can rewrite Equation (4.52) in Lemma 4.16 into

$$s^{-r} \rho_s|_{T^m} = \frac{1}{|\text{Pic}_K^0|} \sum_{\chi_{\ell^*} \in \widehat{T^m}} \rho_{1/s}(\ell^*) \sum_{\chi'|_{T^m} = \chi_{\ell^*}} \chi', \quad (4.54)$$

where we used the identity $|\text{Cl}_K^m| \text{Vol}(T^m) = |\text{Pic}_K^0|$. We will now split up this character decomposition into three parts: the ‘trivial part’, a ‘low-frequency part’ and a ‘high-frequency part’.

The trivial part consists just of the unit character $\mathbf{1} = \mathbf{1}_{\text{Pic}_K^0}$. The low-frequency part consists of those (non-unit) characters $\chi' \in \widehat{\text{Pic}}_K^0$ that are extensions of a character $\chi_{\ell^*} \in \widehat{T^m}$ where $\ell^* \in \Lambda_{K^m,1}$ has a small norm, say, $\|\ell^*\| < r$. Oppositely, the high-frequency part consists of those characters

that are extensions of some $\chi_{\ell^*} \in \hat{T}$ for which $\|\ell^*\| \geq r$. Here, $r \in \mathbb{R}_{>0}$ can in principle be chosen arbitrarily.

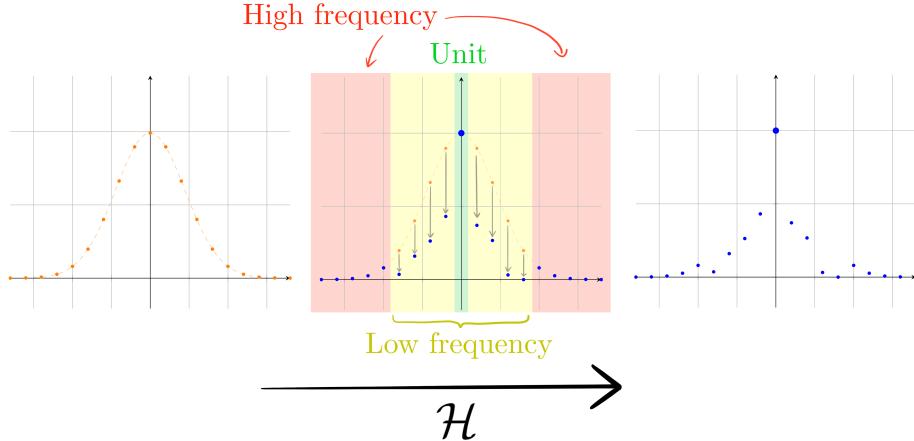


Figure 4.10.: Under assumption of the Extended Riemann Hypothesis, **low-frequency characters** diminish under the Hecke operator, whereas for **high-frequency characters** no such guarantee exists. By taking an initial distribution that has already almost no high-frequency content, like the Gaussian distribution, one can show that applying repeatedly the Hecke operator to such distribution yields an almost-uniform distribution.

$$\begin{aligned}
 |\text{Pic}_{K^m}^0| \cdot s^{-r} \cdot \rho_s |^{T^m} &= \underbrace{\mathbf{1}_{\text{Pic}_{K^m}^0}}_{\text{Unit character}} \\
 &+ \underbrace{\sum_{\substack{\chi_{\ell^*} \in \widehat{T^m} \\ \|\ell^*\| < r}} \rho_{1/s}(\ell^*) \sum_{\substack{\chi' | T^m = \chi_{\ell^*} \\ \chi' \neq \mathbf{1}}} \chi'}_{\text{Low frequency characters}} + \underbrace{\sum_{\substack{\chi_{\ell^*} \in \widehat{T^m} \\ \|\ell^*\| \geq r}} \rho_{1/s}(\ell^*) \sum_{\substack{\chi' | T^m = \chi_{\ell^*}}} \chi'}_{\text{High frequency characters}}, \quad (4.55)
 \end{aligned}$$

Bounding the parts of the character decomposition

Theorem 4.17 (ERH). *Let \mathcal{P} be the set of all prime ideals of a number field K coprime with \mathfrak{m} and with norm at most B , and let $\mathcal{H} = \mathcal{H}_{\mathcal{P}}$ the Hecke operator (see Definition 4.4) for this set of primes. Then, assuming*

the Extended Riemann Hypothesis (see Definition 2.10), and for all $r, s > 0$ with $rs > \sqrt{\frac{r}{4\pi}}$, we have

$$\left\| \mathcal{H}^N(s^{-n} \rho_s|T^m) - \frac{1}{|\text{Pic}_{K^m}^0|} \mathbf{1}_{\text{Pic}_{K^m}^0} \right\|^2 \leq \frac{\rho_{\frac{1}{\sqrt{2}s}}(\Lambda_{K^m,1}^*)}{\text{Vol}(T^m)} \left(c^{2N} + \beta_{\sqrt{2}rs}^{(r)} \right) \quad (4.56)$$

with $c = O\left(\frac{\log(B) \log(B^n \cdot |\Delta_K| \cdot \mathcal{N}(\mathfrak{m}) \cdot (4 + 2\pi r/\sqrt{n})^n)}{B^{1/2}}\right)$.

Proof. The Hecke operator $\mathcal{H} = \mathcal{H}_{\mathcal{P}}$ is a linear operator satisfying $\mathcal{H}(\mathbf{1}_{\text{Pic}_{K^m}^0}) = \mathbf{1}_{\text{Pic}_{K^m}^0}$ and $\mathcal{H}(\chi') = \lambda_{\chi'} \chi'$. Therefore, by applying the Hecke operator N times on Equation (4.55), we obtain

$$\begin{aligned} |\text{Pic}_{K^m}^0| \cdot \mathcal{H}^N(s^{-r} \rho_s|T^m) &= \mathbf{1}_{\text{Pic}_{K^m}^0} + \sum_{\substack{\chi_{\ell^*} \in \widehat{T^m} \\ \|\ell^*\| < r}} \rho_{1/s}(\ell^*) \sum_{\substack{\chi' \mid T^m = \chi_{\ell^*} \\ \chi' \neq \mathbf{1}}} \lambda_{\chi'}^N \chi' \\ &\quad + \sum_{\substack{\chi_{\ell^*} \in \widehat{T^m} \\ \|\ell^*\| \geq r}} \rho_{1/s}(\ell^*) \sum_{\chi' \mid T^m = \chi_{\ell^*}} \lambda_{\chi'}^N \chi', \end{aligned} \quad (4.57)$$

Therefore, by Parseval's theorem (see Equation (2.17)),

$$\begin{aligned} \left\| |\text{Pic}_{K^m}^0| \cdot \mathcal{H}^N(s^{-r} \rho_s|T^m) - \mathbf{1}_{\text{Pic}_{K^m}^0} \right\|^2 &= \underbrace{\sum_{\substack{\chi_{\ell^*} \in \widehat{T^m} \\ \|\ell^*\| < r}} \rho_{1/s}^2(\ell^*) \sum_{\substack{\chi' \mid T^m = \chi_{\ell^*} \\ \chi' \neq \mathbf{1}}} |\lambda_{\chi'}|^{2N}}_{\text{Low frequency}} \\ &\quad + \underbrace{\sum_{\substack{\chi_{\ell^*} \in \widehat{T^m} \\ \|\ell^*\| \geq r}} \rho_{1/s}^2(\ell^*) \sum_{\chi' \mid T^m = \chi_{\ell^*}} |\lambda_{\chi'}|^{2N}}_{\text{High frequency}}. \end{aligned} \quad (4.58)$$

$$\begin{aligned} &+ \underbrace{\sum_{\substack{\chi_{\ell^*} \in \widehat{T^m} \\ \|\ell^*\| \geq r}} \rho_{1/s}^2(\ell^*) \sum_{\chi' \mid T^m = \chi_{\ell^*}} |\lambda_{\chi'}|^{2N}}_{\text{High frequency}}. \end{aligned} \quad (4.59)$$

We will bound the parts Equation (4.58) and Equation (4.59) separately, starting with the share of the latter, the high-frequency characters. By construction, $|\lambda_{\chi'}| \leq 1$ for all $\chi' \in \text{Pic}_{K^m}^0$ (see Lemma 4.5). Combining this with the identity $\rho_{1/s}^2 = \rho_{\frac{1}{\sqrt{2}s}}$ and the fact that there are exactly $|\text{Cl}_K^m|$

character extensions to $\text{Pic}_{K^{\mathfrak{m}}}^0$ for each $\chi_{\ell^*} \in \widehat{T^{\mathfrak{m}}}$ (see [DE16, Cor. 3.6.2]), we have

$$\begin{aligned} \sum_{\substack{\chi_{\ell^*} \in \widehat{T^{\mathfrak{m}}} \\ \|\ell^*\| \geq r}} \rho_{1/s}^2(\ell^*) \sum_{\substack{\chi' \mid T^{\mathfrak{m}} = \chi_{\ell^*}} \mid} |\lambda_{\chi'}|^{2N} &\leq |\text{Cl}_K^{\mathfrak{m}}| \sum_{\substack{\ell^* \in \Lambda_{K^{\mathfrak{m}}, 1}^* \\ \|\ell^*\| \geq r}} \rho_{\frac{1}{\sqrt{2s}}}(\ell^*) \\ &\leq |\text{Cl}_K^{\mathfrak{m}}| \cdot \beta_{\sqrt{2rs}}^{(\mathfrak{r})} \cdot \rho_{\frac{1}{\sqrt{2s}}}(\Lambda_{K^{\mathfrak{m}}, 1}^*), \end{aligned} \quad (4.60)$$

where the last inequality follows from Banaszczyk's tail bound (Lemma 2.25) and the assumption that $rs > \sqrt{\mathfrak{r}/(4\pi)}$.

To bound the share of the low-frequency characters, we need to bound the absolute value of the eigenvalues $\lambda_{\chi'}$ of the low-frequency characters. Invoking the results from analytic number theory in Proposition 4.6 (thus assuming the Extended Riemann Hypothesis) we obtain $|\lambda_{\chi'}| \leq O\left(\frac{\log(B) \log(B^n \cdot |\Delta_K| \cdot \mathcal{N}(\mathfrak{m}) \cdot \mathfrak{q}_{\infty}(\chi'))}{B^{1/2}}\right)$. But since these characters have a 'low frequency', their analytic conductor $\mathfrak{q}_{\infty}(\chi')$ is bounded. More precisely, we have, by Lemma 4.14, that $\mathfrak{q}_{\infty}(\chi') \leq (4 + 2\pi r/\sqrt{n})^n$ for any $\chi' \in \text{Pic}_{K^{\mathfrak{m}}}^0$ such that $\chi' \mid T^{\mathfrak{m}} = \chi_{\ell^*}$ for some $\ell^* \in \Lambda_{K^{\mathfrak{m}}, 1}^*$ with $\|\ell^*\| < r$. Therefore, $|\lambda_{\chi'}| \leq c = O\left(\frac{\log(B) \log(B^n \cdot |\Delta_K| \cdot \mathcal{N}(\mathfrak{m}) \cdot (4 + 2\pi r/\sqrt{n})^n)}{B^{1/2}}\right)$. So, using again the identity $\rho_{1/s}^2 = \rho_{\frac{1}{\sqrt{2s}}}$, and the fact that each χ_{ℓ^*} has $|\text{Cl}_K^{\mathfrak{m}}|$ character extensions to $\text{Pic}_{K^{\mathfrak{m}}}^0$, we have

$$\sum_{\|\ell^*\| \leq r} \rho_{1/s}^2(\ell^*) \underbrace{\sum_{\substack{\chi' \mid T^{\mathfrak{m}} = \chi_{\ell^*}} \mid} |\lambda_{\chi'}|^{2N}}_{\leq |\text{Cl}_K^{\mathfrak{m}}| \cdot c^{2N}} \leq |\text{Cl}_K^{\mathfrak{m}}| \cdot c^{2N} \cdot \rho_{\frac{1}{\sqrt{2s}}}(\Lambda_{K^{\mathfrak{m}}, 1}^*) \quad (4.61)$$

We obtain the result by combining Equations (4.60) and (4.61), dividing by $|\text{Pic}_{K^{\mathfrak{m}}}^0|$ and using the identity $|\text{Pic}_{K^{\mathfrak{m}}}^0| = |\text{Cl}_K^{\mathfrak{m}}| \cdot \text{Vol}(T^{\mathfrak{m}})$. \square

4.3.7. Conclusion

Theorem 4.18. *Let $\varepsilon > 0$ and $s > 0$ be any positive real numbers and let $k \in \mathbb{R}_{>0}$ be a positive real number as well. Let $C \subseteq \Lambda_{K^{\mathfrak{m}}, 1}$ a sub-lattice of the logarithmic ray unit lattice of the number field K . Putting*

$\tilde{s} = \min(\sqrt{2} \cdot s, 1/\eta_1(C^*))$, there exists a bound $B = \tilde{O}(n^{2k}[n^2(\log \log(1/\varepsilon))^2 + n^2(\log(1/\tilde{s}))^2 + n^2 \log([\Lambda_{K^{\mathfrak{m},1}} : C])^2 + (\log(|\Delta_K| \mathcal{N}(\mathfrak{m})))^2])$ such that for any integer

$$N \geq \left\lceil \frac{1}{2k \log n} \cdot (\mathfrak{r} \cdot \log(1/\tilde{s}) + \log|\text{Pic}_{K^{\mathfrak{m}}}^0| + 2 \log(1/\varepsilon) + \log[\Lambda_{K^{\mathfrak{m},1}} : C] + 2) \right\rceil, \quad (4.62)$$

the random walk distribution $\mathcal{W}_{\text{Pic}_{K^{\mathfrak{m}}}^0}(B, N, s)$ is ε -close to uniform in $L_1(\text{Pic}_{K^{\mathfrak{m}}}^0)$, i.e.,

$$\left\| \mathcal{W}_{\text{Pic}_{K^{\mathfrak{m}}}^0}(B, N, s) - \mathcal{U}(\text{Pic}_{K^{\mathfrak{m}}}^0) \right\|_1 \leq \varepsilon.$$

Proof. Let $1 > \varepsilon > 0$, $s > 0$ and $k \in \mathbb{R}_{>0}$ be given. Let $C \subseteq \Lambda_{K^{\mathfrak{m},1}} = \text{Log}(\mathcal{O}_K^\times \cap K^{\mathfrak{m},1})$ be a sublattice of the logarithmic ray unit lattice of index $[\Lambda_{K^{\mathfrak{m},1}} : C]$. Since, by construction, $1/\tilde{s} \geq \eta_1(C^*)$ and $1/\tilde{s} \geq 1/(\sqrt{2}s)$, and since $\Lambda_{K^{\mathfrak{m},1}}^* \subseteq C^*$, we have

$$\begin{aligned} \rho_{\frac{1}{\sqrt{2}s}}(\Lambda_{K^{\mathfrak{m},1}}^*) &\leq \rho_{\frac{1}{\sqrt{2}s}}(C^*) \leq \rho_{1/\tilde{s}}(C^*) \leq 2 \cdot \det(C)/\tilde{s}^{\mathfrak{r}} \\ &\leq 2 \cdot [\Lambda_{K^{\mathfrak{m},1}} : C] \cdot \text{Vol}(T^{\mathfrak{m}})/\tilde{s}^{\mathfrak{r}} \end{aligned} \quad (4.63)$$

Using this inequality and Hölder's inequality (i.e., $\|f \cdot 1\|_1 \leq \|f\|_2 \|1\|_2$), noting that $\|1_{\text{Pic}_{K^{\mathfrak{m}}}^0}\|_2^2 = |\text{Pic}_{K^{\mathfrak{m}}}^0|$ and applying Theorem 4.17 and Equation (4.63), we obtain, for each $r > \sqrt{\mathfrak{r}}/(\sqrt{2}s)$,

$$\begin{aligned} &\left\| \mathcal{W}_{\text{Pic}_{K^{\mathfrak{m}}}^0}(B, N, s) - \mathcal{U}(\text{Pic}_{K^{\mathfrak{m}}}^0) \right\|_1^2 \\ &\leq |\text{Pic}_{K^{\mathfrak{m}}}^0| \cdot \left\| \mathcal{H}^N(s^{-\mathfrak{r}} \rho_s|T^{\mathfrak{m}}|) - \frac{1}{|\text{Pic}_{K^{\mathfrak{m}}}^0|} \mathbf{1}_{\text{Pic}_{K^{\mathfrak{m}}}^0} \right\|_2^2 \\ &\leq |\text{Cl}_K^{\mathfrak{m}}| \cdot \rho_{\frac{1}{\sqrt{2}s}}(\Lambda_{K^{\mathfrak{m},1}}^*) \cdot (c^{2N} + \beta_{\sqrt{2}rs}^{(\mathfrak{r})}) \\ &\leq |\text{Cl}_K^{\mathfrak{m}}| \cdot 2 \cdot \text{Vol}(T^{\mathfrak{m}}) \cdot [\Lambda_{K^{\mathfrak{m},1}} : C] \cdot \tilde{s}^{-\mathfrak{r}} \cdot (c^{2N} + \beta_{\sqrt{2}rs}^{(\mathfrak{r})}) \\ &\leq 2 \cdot |\text{Pic}_{K^{\mathfrak{m}}}^0| \cdot [\Lambda_{K^{\mathfrak{m},1}} : C] \cdot \tilde{s}^{-\mathfrak{r}} \cdot (c^{2N} + \beta_{\sqrt{2}rs}^{(\mathfrak{r})}). \end{aligned} \quad (4.64)$$

Here, $c = O\left(\frac{\log(B) \log(B^n \cdot |\Delta_K| \cdot \mathcal{N}(\mathfrak{m}) \cdot (4+2\pi r/\sqrt{n})^n)}{B^{1/2}}\right)$, as in Theorem 4.17. We proceed by bounding the two summands in Equation (4.64) separately.

- By putting³ r equal to the maximum of $\sqrt{r}/(\sqrt{2}s)$ and

$$\frac{1}{\sqrt{2}s} \cdot \sqrt{2 + r \log(1/\tilde{s}) + 2 \log(1/\varepsilon) + \log|\text{Pic}_{K^m}^0| + \log[\Lambda_{K^m,1} : C]}$$

we deduce that $2 \cdot |\text{Pic}_{K^m}^0| \cdot [\Lambda_{K^m,1} : C] \cdot \tilde{s}^{-r} \cdot \beta_{\sqrt{2}rs}^{(r)} \leq \varepsilon^2/2$.

- Subsequently, choose⁴ a $B = \tilde{O}(n^{2k}[\log(|\Delta_K| \mathcal{N}(\mathfrak{m}))^2 + n^2 \log(r)^2])$, i.e.,

$$\begin{aligned} B = \tilde{O} & \left(n^{2k} \cdot [\log(|\Delta_K| \mathcal{N}(\mathfrak{m}))^2 + n^2 \log(1/\tilde{s})^2 \right. \\ & \left. + n^2 \log(\log(1/\varepsilon))^2 + n^2 \log([\Lambda_{K^m,1} : C])^2] \right) \end{aligned}$$

such that $c \leq 1/n^k$, where $c = O\left(\frac{\log(B) \log(B^n \cdot |\Delta_K| \mathcal{N}(\mathfrak{m}) \cdot (4 + 2\pi r/\sqrt{n})^n)}{B^{1/2}}\right)$ as in Theorem 4.17. Finally, taking any integer $N \geq \frac{1}{k \log n} \cdot \left(\frac{r}{2} \cdot \log(1/\tilde{s}) + 2 \log(1/\varepsilon) + \frac{1}{2} \log|\text{Pic}_{K^m}^0| + \log[\Lambda_{K^m,1} : C] + 1\right)$ and noting that $c^{\frac{1}{k \log n}} \leq 1/e$, we deduce that $2 \cdot |\text{Pic}_{K^m}^0| \cdot [\Lambda_{K^m,1} : C] \cdot \tilde{s}^{-r} \cdot c^{2N} \leq \varepsilon^2/2$.

Combining, we can bound the right-hand side of Equation (4.64) by ε^2 . Taking square roots gives the final result. \square

Remark 4.19. Consider the base case $\mathfrak{m} = \mathcal{O}_K$. The occurrence of the sublattice $C \subseteq \Lambda_K$ of the log unit lattice in Theorem 4.18 might appear strange at first sight — indeed, just taking $C = \Lambda_K$ would make the result less complex and seemingly about equally powerful.

We chose to phrase Theorem 4.18 in this way, because, in some number fields, certain subgroups of the unit group are better understood than the full unit group itself. For cyclotomic number fields, for example, the structure of the subgroup of the cyclotomic units is simpler than that of the full unit group. Due to this simpler structure, we can achieve a tighter bound on $\eta_1(C^*)^r \cdot [\Lambda_K : C]$ than we have on $\eta_1(\Lambda_K)$ (where C is here chosen to be the logarithmic image of the cyclotomic units).

³We use the bound $\beta_\alpha^{(r)} \leq e^{-\alpha^2}$ for $\alpha \geq \sqrt{r}$

⁴In this bound on B one would expect an additional $\log|\text{Pic}_{K^m}^0|$. But as it is bounded by $\log(\log(|\Delta_K| \mathcal{N}(\mathfrak{m})))$ (see Lemma 2.17), it can be put in the hidden polylogarithmic factors.

Such a tight upper bound on the product $\eta_1(C^*)^r \cdot [\Lambda_K : C]$ is important. This product does namely not only have a large influence on the complexity, but also has a significant leverage on $B^{N/n}$, the quality loss in the shortest-vector problem of the reduction in Chapter 5.

By taking $C = \Lambda_{K^m}$ in Theorem 4.18, we obtain the main theorem.

Theorem 4.3 (Random Walks on the Arakelov Ray Class Group, ERH). *Let $\varepsilon > 0$ and $s > 0$ be any positive real numbers and let $k \in \mathbb{R}_{>0}$ be a positive real number as well. Putting $\tilde{s} = \min(\sqrt{2} \cdot s, 1/\eta_1(\Lambda_{K^m,1}^*))$, there exists a bound $B = \tilde{O}(n^{2k}[n^2(\log \log(1/\varepsilon))^2 + n^2(\log(1/\tilde{s}))^2 + (\log(|\Delta_K| \mathcal{N}(\mathfrak{m})))^2])$ such that for any integer*

$$N \geq \left\lceil \frac{1}{2k \log n} \cdot (\mathfrak{r} \cdot \log(1/\tilde{s}) + \log |\text{Pic}_{K^m}^0| + 2 \log(1/\varepsilon) + 2) \right\rceil, \quad (4.46)$$

the random walk distribution $\mathcal{W}_{\text{Pic}_{K^m}^0}(B, N, s)$ is ε -close to uniform in $L_1(\text{Pic}_{K^m}^0)$, i.e.,

$$\left\| \mathcal{W}_{\text{Pic}_{K^m}^0}(B, N, s) - \mathcal{U}(\text{Pic}_{K^m}^0) \right\|_1 \leq \varepsilon.$$

Remark 4.20. Consider again the base case $\mathfrak{m} = \mathcal{O}_K$. The running time of the random walk as in Theorem 4.18 depends on quite a subtle way on the Gaussian spread s of the continuous walk. Roughly said, one can distinguish three regions; $s < 1/\eta_1(\Lambda_K^*)$, $1/\eta_1(\Lambda_K^*) \leq s < \eta_1(\Lambda_K)$ and $\eta_1(\Lambda_K) \leq s$, see also Figure 4.11.

- (i) If $s < 1/\eta_1(\Lambda_K^*)$, the Gaussian is narrow compared to the unit group torus T . Each Gaussian covers a volume of around $s^{\mathfrak{r}}$ in the Arakelov class group, which has volume $|\text{Pic}_K^0|$. It is then intuitively clear that around $O(|\text{Pic}_K^0|/s^{\mathfrak{r}})$ reasonably equidistributed duplicates of that Gaussian spot are needed to cover the entire Arakelov class group, i.e., to get a nearly uniform distribution. As the duplicates grow exponentially per random walk step, one expects to need $O(\log |\text{Pic}_K^0| + \mathfrak{r} \log(s^{-1}))$ random walk steps. So in this particular case, the inverse $1/s$ of the Gaussian spread s has a significant influence on the running time.

- (ii) If $1/\eta_1(\Lambda_K^*) \leq s < \eta_1(\Lambda_K)$, the Gaussian spread is already so large that it has already some overlap on the unit torus. So it is then to be expected that the running time of the random walk does not so much depend on this Gaussian spread *per se*, but rather on the structure of the log unit lattice. If there is a significant gap between $1/\eta_1(\Lambda_K^*) \approx 1/\lambda_r(\Lambda_K^*) \approx \lambda_1(\Lambda_K)$ and $\eta_1(\Lambda_K) \approx \lambda_r(\Lambda_K)$, one can deduce that this log unit lattice must be quite ‘distorted’.
- (iii) If $\eta_1(\Lambda_K) \leq s$, the Gaussian is already so wide that it covers the entire unit group torus (the connected component of the unit in Pic_K^0). Then neither the Gaussian spread s nor the log unit lattice Λ_K have then any influence on the running time: For such large s one can simply replace $\log(1/\tilde{s})$ by 0 and $\text{Vol}(\text{Pic}_K^0)$ by h_K in Theorem 4.18. Not surprisingly, one then recovers the ‘rapid mixing theorem’ for ideal class groups by Jetchev and Wesolowski [JW15]. Additionally, by letting k tend to zero (i.e., allowing for an infinite number of steps N) one obtains that the prime ideals of norm below $B = \tilde{O}(\log(|\Delta_K|)^2)$ generate the ideal class group, a fact better known as Bach’s bound [Bac90].

For the case $\mathfrak{m} \neq \mathcal{O}_K$, the same reasoning applies, but with the ray unit torus $T^{\mathfrak{m}}$ instead.

Applications of the Random Walk theorem in the subsequent chapters

The genericness of the random walk theorem (Theorem 4.18) allows it to be used for many applications. In this thesis, we specialize the parameters for two cases.

The first case concerns the worst-case to average-case reduction for Hermite-SVP on ideal lattices, the topic of Chapter 5. In that chapter we apply Theorem 4.18 for general number fields and cyclotomic fields separately (see Proposition 5.10). In this specialization of the random walk theorem we aim at an as low as possible value for $B^{N/n}$, as this is the loss in shortness quality in the worst-case to average-case reduction. The cyclotomic field gets a special treatment because one can obtain sharper bounds in that case,

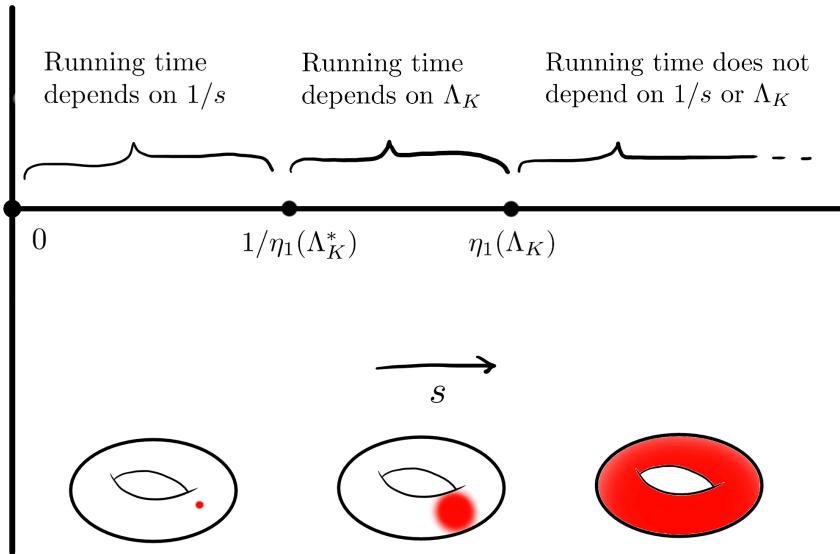


Figure 4.11.: The larger the Gaussian deviation s is, the less influence it has on the running time of the random walk.

assuming that the class number of the maximal totally real subfield of the cyclotomic field remains reasonably small, asymptotically. This is due to the occurrence of the *cyclotomic units* in cyclotomic fields, a subgroup of the unit group with particularly nice properties.

The second case concerns ideal sampling, the main topic of Chapter 6. In that chapter we show that a specific way of sampling in an ideal resembles sampling in a ideal that is a result of a random walk on the Arakelov class group. It applies Theorem 4.18 to show that this latter result of a random walk is close to uniformly distributed. In those cases one can apply ideal density results to lower bound the probability that the relative ideal (the sampled element divided by the input ideal) lies in a certain ideal set. In this application there is not really a restriction on any parameter, apart that those need to be small enough for the sampling algorithm to be efficient. In Theorem 6.3 of Chapter 6 we made some parameter choices (e.g., $s \approx 1/n^2$) to make the computation less involved and making the sampling in the

distorted box feasible; also the specialized theorem in this chapter is slightly aimed at a small B , the bound on the norm of the prime ideals. For the ideal sampling result this is not required; one can append these parameters as one wishes, provided that the sampling in the distorted box can still be done efficiently (or is at least non-vacuous) and taking care that the prime ideals involved (this depends on the parameter B) do not get too large to be feasible.

5. A Worst-case to Average-case Reduction for Ideal Lattices

5.1. Summary

In this chapter we achieve a worst-case to average-case reduction for the Hermite Shortest Vector Problem (SVP) on ideal lattices of a fixed number field. Such a reduction allows to transform a fixed chosen instance of a problem (the worst case) to a sample of a fixed *distribution* over all instances of this problem (the average case). Slightly more formally said, a worst-case to average-case reduction consists of two parts: the first one being a *definition of the average-case distribution* and the second one being an *algorithm that reduces any input instance to a sample of that average-case distribution*.

In the reduction of this chapter, which concerns Hermite-SVP on ideal lattices of a fixed number field, this average-case distribution will be defined as something closely related to the *uniform distribution on the Arakelov class group*. This Arakelov class group is essentially the group of ideal lattices up to isometry.

The reduction algorithm in this chapter transforms any fixed input ideal lattice to a sample of the average-case distribution on the Arakelov class group by means of a *random walk*, as introduced in the previous chapter. This ‘random walk’ transformation of the input ideal lattice only slightly changes its geometry and is therefore compatible with the Hermite Shortest Vector Problem. More concretely, any short vector of the transformed ideal lattice can be reasonably *untransformed* to yield a short vector of the input lattice, with only a small loss in quality.

This particular approach for a worst-case to average-case reduction faces two challenges. The first challenge consists of finding a suitable *representation* of ideal lattices (or Arakelov classes), whereas the second one involves an appropriate treatment of the inherently continuous ideal lattices on finite precision machines.

Such a representation of ideal lattices suitable for the purposes of the worst-case to average-case reduction turns out to be doable by means of a *distribution* over the group of fractional ideals. More precisely, with any fixed ideal lattice we associate an algorithm that efficiently samples from a specific distribution, mainly consisting of fractional ideals that geometrically resemble the input ideal lattice – i.e., whose Arakelov class is close to that of the original ideal lattice. This specific distribution is then our representation of that fixed ideal lattice.

The appropriate treatment of the inherently continuous objects on finite machines happens by *discretization*. A considerable amount of this chapter is devoted to showing that this discretization does not have a significant effect on the overall worst-case to average-case reduction.

5.2. Introduction

The space of all ideal lattices (up to isometry) in a given number field forms naturally an abelian group, called the *Arakelov class group* – a fact well known to number theorists (e.g., [Sch08]). Yet this notion has never appeared explicitly in the literature on lattice-based cryptography. The relevance of this perspective is already illustrated by some previous work which implicitly exploit Arakelov ideals [Eis+14; BS16] and even the Arakelov class group [PHS19; Lee+19]. Beyond its direct result, this chapter aims at highlighting this powerful Arakelov class group formalism for finer and more rigorous analysis of computational problems in ideal lattices.

5.2.1. The Result

We exploit the random walk theorem of Chapter 4 to relate the average-case and the worst-case of Ideal-SVP, due to the interpretation of the Arakelov class group as the space of all ideal lattices up to isometry. Note that this reduction does not directly impact the security of existing schemes: there exists no modern cryptographic scheme based on the average-case version of Ideal-SVP. The value of our result lies in the introduction of a new tool, and an illustration of the cryptanalytic insights it offers.

As already mentioned, ideal lattices (up to isometry) of a given number field K can be identified with the elements of the Arakelov class group, also known as the degree zero part Pic_K^0 of the Picard group. There are two ways to move within this group: given an ideal, one can obtain a new one by ‘distorting’ it, or by ‘sparsifying’ it. In both cases, finding a short vector in the target ideal also allows to find a short vector in the source ideal, up to a certain loss of shortness. So, the quality (i.e., the shortness) of the short vector deteriorates with each extra step of the walk; therefore, we minimize the length of the random walk subject to the requirement that the target ideal is uniformly randomly distributed in the Arakelov class group.

This approach leads to a surprisingly tight reduction. In the case of cyclotomic number fields of prime power conductor $m = p^k$, under the Riemann Hypothesis for Hecke L -functions (which we abbreviate ERH for the Extended Riemann Hypothesis), and a mild assumption on the structure of the class groups, the loss of approximation factor is as small as $\tilde{O}(\sqrt{m})$. In other words:

Main Theorem (informal). *Let $m = p^k$ be a prime power. If there exists a polynomial-time algorithm for solving Hermite-SVP with approximation factor γ over random ideal lattices of $\mathbb{Q}(\zeta_m)$, then there also exists a polynomial time algorithm that solves Hermite-SVP in any ideal lattice with approximation factor $\gamma' = \gamma \cdot \sqrt{m} \cdot \text{poly}(\log m)$.*

In fact, this theorem generalizes to all number fields, but the loss in approximation factor needs to be expressed in more involved quantities. The precise

statement is the object of Theorem 5.9.

5.2.2. Overview

The Arakelov class group. Both the unit group [Cra+16] and the class group [CDW17] have been shown to play a key role in the cryptanalysis of ideal lattice problems. In these works of Cramer et al. [Cra+16; CDW17], these groups are exploited independently, in ways that nevertheless share strong similarities with each other. More recently, both groups have been used in combination for cryptanalytic purposes [PHS19; Lee+19]. It therefore seems natural to turn to a unifying theory.

The Arakelov class group (denoted Pic_K^0) is a combination of the unit torus $T = \text{Log } K_{\mathbb{R}}^0 / \text{Log}(\mathcal{O}_K^\times)$ and of the class group Cl_K . The exponent 0 in $K_{\mathbb{R}}^0$ refers to elements of algebraic norm 1 (i.e., modulo renormalization), while the subscript \mathbb{R} indicates that we are working in the topological completion of K . By ‘a combination’ we do not exactly mean that Pic_K^0 is a direct product; we mean that there is a short exact sequence

$$0 \longrightarrow T \longrightarrow \text{Pic}_K^0 \longrightarrow \text{Cl}_K \longrightarrow 0.$$

That is, T is (isomorphic to) a subgroup of Pic_K^0 , and Cl_K is (isomorphic to) the quotient Pic_K^0 / T . The Arakelov class group is an abelian group which combines an uncountable (yet compact) part T and a finite part Cl_K ; topologically, it should be thought of as $|\text{Cl}_K|$ many disconnected copies of the torus T (see Figure 4.1).

A worst-case to average-case reduction for ideal-SVP. An important aspect of the Arakelov class group for the present work is that this group has a geometric interpretation: it can essentially be understood as the group of all ideal lattices up to K -linear isometries. Furthermore, being equipped with a metric, it naturally induces a notion of near-isometry. Such a notion gives a new handle to elucidate the question of the hardness of ideal-SVP. Namely, knowing a short vector in \mathfrak{a} , and a near-isometry from \mathfrak{a} to $\tilde{\mathfrak{a}}$, one can

deduce a short vector of $\tilde{\mathfrak{a}}$ up to a small loss induced by the distortion of the near-isometry. This suggests a strategy towards a worst-case to average-case reduction for ideal lattices, namely by randomly distorting a worst-case ideal to a random one (see Figure 5.2).

However, there are two issues with this strategy: first, this near-isometry keeps staying in a fixed class of Cl_K ; i.e., one is stuck in one of the potentially many separated copies of the torus that constitute the Arakelov class group. Second, even if $|\text{Cl}_K| = 1$, the unit torus T might be too large, and to reach the full torus from a given point, one may need near-isometry that are too distorted for our purposes.

In the language of algebraic geometry, distortion of ideal lattices corresponds to the ‘infinite places’ of the field K , while we can also exploit the ‘finite places’, i.e., the prime ideals. Indeed, if \mathfrak{c} is an integral ideal of small norm and $\tilde{\mathfrak{a}} = \mathfrak{c}\mathfrak{a}$, then $\tilde{\mathfrak{a}}$ is a sublattice of \mathfrak{a} and a short vector of $\tilde{\mathfrak{a}}$ is also a somewhat short vector of \mathfrak{a} , an idea already used in [CDW17; PHS19] (see Figure 5.1).

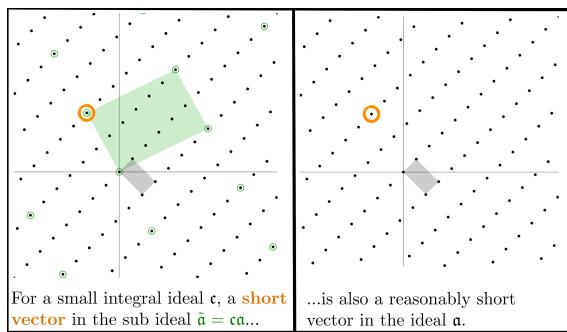


Figure 5.1.: If \mathfrak{c} is an integral ideal of small norm and $\tilde{\mathfrak{a}} = \mathfrak{c}\mathfrak{a}$, then $\tilde{\mathfrak{a}}$ is a sublattice of \mathfrak{a} and a short vector of $\tilde{\mathfrak{a}}$ is also a somewhat short vector of \mathfrak{a} .

5.2.3. Related work

Relation to recent cryptanalytic works. The general approach to this result was triggered by a heuristic observation made in [DPW19], suggesting that

the worst-case behavior of the quantum ideal-SVP algorithm built out of [Eis+14; BS16; Cra+16; CDW17] could be made not that far of the average-case behavior they studied experimentally. More specifically, we do achieve the hoped generalization of the class group mixing theorem of [JMV09; JW15] to Arakelov class groups.

Prior self-reduction via random walks. As already mentioned, our result shares strong similarities with a technique introduced by Jao, Miller and Venkatesan [JMV09] to study the discrete logarithm problem on elliptic curves. Just as ideal lattices can be seen as elements of the Arakelov class group, elliptic curves in certain families are in bijective correspondence with elements of the class group of a quadratic imaginary number field. In [JMV09], Jao et al. studied (discrete) random walks on class groups, and showed that they have a rapid mixing property. They deduced that from any elliptic curve, one can efficiently construct a random isogeny (a group homomorphism) to a uniformly random elliptic curve, allowing to transfer a worst case instance of the discrete logarithm problem to an average case instance. Instead of the finite class group, we studied random walks on the infinite Arakelov class group, which led to consequences in lattice-base cryptography, an area seemingly unrelated to elliptic curve cryptography.

Prior self-reduction for ideal lattices. Our self-reducibility result is not the first of its kind: in 2010, Gentry already proposed a self-reduction for an ideal lattice problem [Gen10], as part of his effort of basing Fully-Homomorphic Encryption on worst-case problems [Gen09]. Our result differs in several points.

- Our reduction does not rely on a factoring oracle, and is therefore classically efficient; this was already advertised as an open problem in [Gen10].
- The reduction of Gentry considers the Bounded Distance Decoding problem (BDD) in ideal lattices rather than a short vector problem. Note that this distinction is not significant with respect to quantum computers [Reg09].

- The definition of average case distribution is significantly different, and we view the one of [Gen10] as being somewhat ad-hoc. Given that the Arakelov class group captures exactly ideal lattices up to isometry, we consider the uniform distribution in the Arakelov class group as a much more natural and conceptually simpler choice.
- The worst case ideal input of [Gen10] has restrictions on the size of the norm, whereas our worst case ideal input is unrestricted.
- The loss on the approximation factor of our reduction is much more favorable than the one of Gentry [Gen10]. For example, in the case of cyclotomic number fields with prime-power conductor, Gentry's reduction (on BDD) seems to loose a factor at least $\Theta(n^{4.5})$, while our reduction (on Hermite-SVP) only loses a factor $\tilde{O}(\sqrt{n})$ making a mild assumption on plus-part h^+ of the class number.

Structure of this chapter

We start the remainder of this chapter by constructing an representation of Arakelov class elements that is appropriate to use in a worst-case to average-case reduction (Section 5.3).

After that, we describe a simplified version of the worst-case to average-case reduction; we leave out the difficulties concerning finite machine precision (Section 5.4). In the last part of this chapter, we will show by quite technical means that ignoring finite precision does not impact the reduction significantly (Section 5.5).

5.3. Representation of Ideal Lattices by Means of Distributions

Ideal lattices

Though the notion of ideal lattices is already given in this thesis (see Definition 2.19), we will restate the definition here.

Definition 5.1 (Ideal lattices). *Let K be a number field with ring of integers \mathcal{O}_K . An ideal lattice of K is a \mathcal{O}_K -module $I \subseteq K_{\mathbb{R}}$, with the additional requirement that there exists an $x \in K_{\mathbb{R}} \setminus \{0\}$ such that $xI \subseteq \mathcal{O}_K$. We denote the group of ideal lattices by IdLat_K .*

In essence, the group of ideal lattices IdLat_K can be considered as a sort-of completion of the group of fractional ideals \mathcal{I}_K , in the same sense that the reals \mathbb{R} are a completion of \mathbb{Q} . A straightforward way to imagine an ideal lattice $x\mathfrak{a} \subseteq K_{\mathbb{R}}$ is to think of an ideal $\mathfrak{a} \subseteq \mathcal{O}_K$ that is ‘perturbed’ by a vector $x \in K_{\mathbb{R}} = \{y \in \bigoplus_{\sigma: K \rightarrow \mathbb{C}} \mathbb{C} \mid y_{\bar{\sigma}} = \overline{y_{\sigma}}\}$ (see Figure 5.2).

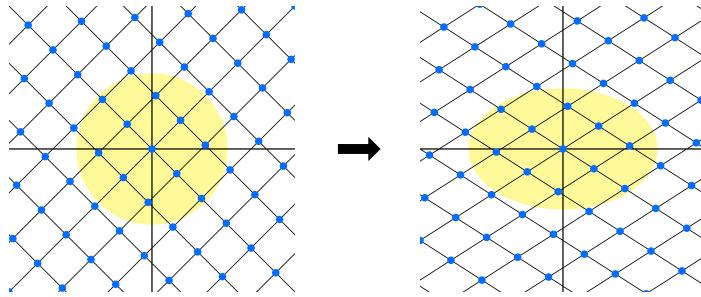


Figure 5.2.: In this two-dimensional example, the left ideal lattice is slightly stretched in the x -direction and slightly shrunk in the y -direction, leading to the perturbed ideal lattice on the right. The yellow circle functions as a visual aid, making the precise deformation of the lattice more explicit.

Representations

Above interpretation immediately gives a representation of the ideal lattice $x\mathfrak{a}$ by the pair $(x, \mathfrak{a}) \in K_{\mathbb{R}} \times \mathcal{I}_K$. But that representation is by no means unique; indeed, one can check that $(x\alpha^{-1}, (\alpha)\mathfrak{a}) \in K_{\mathbb{R}} \times \mathcal{I}_K$, for example, generates the same ideal lattice for any $\alpha \in K^*$. Here, $\alpha^{-1} \in K$ is seen as an element in $K_{\mathbb{R}}$ via the Minkowski embedding $K \hookrightarrow K_{\mathbb{R}}$.

Why do we need an efficient and canonical representation of an ideal lattice (or Arakelov class)?

As mentioned before, a worst-case to average-case distribution of a certain set of problem instances consists of two parts: the definition of a distribution on this set of instances, and an algorithm that reduces any fixed problem instance to this distribution.

Given an Arakelov divisor $\mathbf{a} \in \text{Div}_K^0$, we know how to randomize it so that it is uniformly random in the quotient group Pic_K^0 ; namely, by the random walk procedure (see Chapter 4). So, for any $\mathbf{a} \in \text{Div}_K^0$ we can efficiently compute a distribution $\mathbb{D}_{\mathbf{a}} \in L_1(\text{Div}_K^0)$ that becomes an uniform distribution under the canonical map $L_1(\text{Div}_K^0) \rightarrow L_1(\text{Pic}_K^0)$, $\mathbb{D} \mapsto \sum_{k \in K^*/\mu_K} \mathbb{D}(\cdot + k)$.

To obtain a worst-case to average-case reduction we need a *fixed* (average-case) distribution \mathbb{D}_0 on Div_K^0 and an efficient (reduction) map

$$\psi : L_1(\text{Div}_K^0) \rightarrow L_1(\text{Div}_K^0)$$

such that for all $\mathbf{a} \in \text{Div}_K^0$, $\psi(\mathbb{D}_{\mathbf{a}}) = \mathbb{D}_0$. Also, this reduction map must be preserving certain geometric properties (be Hermite-SVP compatible) to be an actual useful reduction map.

Suppose for the moment that one has a canonical ‘lift’ $\mathbb{L} : \text{Div}_K^0 \rightarrow \text{Div}_K^0$ for which holds $[\mathbf{a}] = [\mathbf{b}] \Rightarrow \mathbb{L}(\mathbf{a}) = \mathbb{L}(\mathbf{b})$; i.e., it ‘factors through’ Pic_K^0 . And suppose that this map is compatible with Hermite-SVP, i.e., solving Hermite-SVP in $\mathbb{L}(\mathbf{a})$ allows to solve Hermite-SVP in \mathbf{a} . Then this lift \mathbb{L} serves as a reduction map, by sending the distribution $\mathbb{D}_{\mathbf{a}} \in L_1(\text{Div}_K^0)$ to $\mathbb{L}(\mathbb{D}_{\mathbf{a}})$, with which we mean the distribution that samples $\mathbb{L}(\mathbf{b})$ with density $\mathbb{D}_{\mathbf{a}}(\mathbf{b})$. By the fact that $\mathbb{D}_{\mathbf{a}}$ maps to the uniform distribution under the canonical map $L_1(\text{Div}_K^0) \rightarrow L_1(\text{Pic}_K^0)$, the distribution $\mathbb{L}(\mathbb{D}_{\mathbf{a}}) = \mathbb{D}_0$ is the same for all $\mathbf{a} \in \text{Div}_K^0$. So, such an efficient and Hermite-SVP compatible lift \mathbb{L} , which then computes an efficient and canonical representation of ideal lattices, is essentially what remains to construct to make the worst-case to average-case reduction work.

Algorithm 3: Sampling from the distribution $\mathcal{D}_{x\mathfrak{a}}$.

Require: A pair $(x, \mathfrak{a}) \in K_{\mathbb{R}} \times \mathcal{I}_K$ such that $\mathcal{N}(\mathfrak{a}) \prod_{\sigma} x_{\sigma} = 1$.

Ensure:

- A sample \mathfrak{d}^{-1} from the distribution $\mathcal{D}_{x\mathfrak{a}}$
- A $v \in x\mathfrak{a}$ such that $\mathfrak{d}^{-1} = v^{-1}x\mathfrak{a}$.

- 1: Put $\varsigma = 2^{n+1} \cdot n \cdot |\Delta_K|$ and $M = 2\sqrt{n} \cdot \varsigma$.
- 2: Sample a center $c = (c_{\sigma})_{\sigma}$ uniformly in $\mathcal{C}_M = \{(y_{\sigma})_{\sigma} \mid |y_{\sigma}| = M \text{ for all embeddings } \sigma\}$.
- 3: Sample from the discrete Gaussian $\mathcal{G}_{x\mathfrak{a}, \varsigma, c}$ with respect to the ideal lattice $x\mathfrak{a}$ with center $c = (c_{\sigma})_{\sigma}$ and standard deviation ς , leading to some $v \in x\mathfrak{a}$.
- 4: **return** the inverse integral ideal $\mathfrak{d}^{-1} = v^{-1}x\mathfrak{a} \in \mathcal{I}_K$ and $v \in x\mathfrak{a}$.

We could not find such an efficient map \mathbb{L} that is also compatible with Hermite-SVP – instead, we use a map $\mathbb{L} : \text{Div}_K^0 \rightarrow L_1(\text{Div}_K^0)$ that is sufficient for our needs. This is a canonical representation by means of a *distribution*. The map we chose has even codomain $L_1(\mathcal{I}_K)$, i.e., involves a discretization for efficiency. So the map $\mathbb{L} : \text{Div}_K^0 \rightarrow L_1(\mathcal{I}_K)$ we construct, satisfies $\mathbb{L}(\mathbf{a}) = \mathbb{L}(\mathbf{b})$ for $[\mathbf{a}] = [\mathbf{b}]$ and is compatible with Hermite-SVP.

Concretely, $\mathbb{L}(\mathbf{a})$ consists of sampling a ‘balanced’ element $\alpha \in \mathbf{a}$ and outputting the ideal $\alpha^{-1} \cdot \mathbf{a} \in \mathcal{I}_K$. This ideal then quite resembles the geometry of \mathbf{a} and lies in the same ideal class; so this ideal (when reduced to the Arakelov class group) must be close to $[\mathbf{a}]$.

Representation by means of a distribution

A representation that is both unique and (in some sense) classically efficiently computable can be made by means of a *distribution*. We will define a map $\text{IdLat}_K \rightarrow L_1(\mathcal{I}_K)$, $x\mathfrak{a} \mapsto \mathcal{D}_{x\mathfrak{a}}$ having the property that $\mathcal{D}_{x\mathfrak{a}}$ is an efficiently samplable distribution for any input ideal lattice $x\mathfrak{a} \in \text{IdLat}_K$. The computation of this map $x\mathfrak{a} \mapsto \mathcal{D}_{x\mathfrak{a}}$ is described in Algorithm 3.

Remark 5.2. It follows from the description of Algorithm 3 that the distribution $\mathcal{D}_{x\mathfrak{a}}$ indeed depends only on the ideal lattice $x\mathfrak{a}$ and not so much on the representation $(x, \mathfrak{a}) \in K_{\mathbb{R}} \times \mathcal{I}_K$ thereof. Hence the notation $\mathcal{D}_{x\mathfrak{a}}$, instead of, for example, $\mathcal{D}_{(x, \mathfrak{a})}$.

The output of the element $v \in x\mathfrak{a}$ such that $\mathfrak{d}^{-1} = v^{-1}x\mathfrak{a}$ does not take any part in the distribution $\mathcal{D}_{x\mathfrak{a}}$. But it will have a major role in the worst-case to average-case reduction (see Algorithm 4), because it relates \mathfrak{d}^{-1} to the input ideal lattice $x\mathfrak{a}$.

Equivalently, the distribution $\mathcal{D}_{x\mathfrak{a}}$ can be described by the following definition.

Definition 5.3 (Distribution representation of ideal lattices). Let $(x, \mathfrak{a}) \in K_{\mathbb{R}} \times \mathcal{I}_K$. The distribution $\mathcal{D}_{x\mathfrak{a}} \in L_1(\mathcal{I}_K)$ is supported only by inverse integral ideals. For integral ideals $\mathfrak{d} \in \mathcal{I}_K$ the probability is defined by the following rule.

$$\mathcal{D}_{x\mathfrak{a}}[\mathfrak{d}^{-1}] = \frac{1}{\text{Vol}(\mathcal{C}_M)} \int_{c \in \mathcal{C}_M} \frac{1}{\rho_{\varsigma}(x\mathfrak{a} - c)} \sum_{\substack{v \in x\mathfrak{a} \\ (v) = x\mathfrak{a}\mathfrak{d}}} \rho_{\varsigma}(v - c) dc, \quad (5.65)$$

where $\varsigma = 2^{n+1} \cdot n \cdot |\Delta_K|$, $M = 2\sqrt{n} \cdot \varsigma$ and $\mathcal{C}_M = \{(x_{\sigma})_{\sigma} \in K_{\mathbb{R}} \mid |x_{\sigma}| = M\}$, the M -hypersphere in $K_{\mathbb{R}}$.

The fact that $\mathcal{D}_{x\mathfrak{a}}$ is a distribution follows by the following computation.

$$\begin{aligned} \sum_{\mathfrak{d} \in \mathcal{I}_K} \mathcal{D}_{x\mathfrak{a}}[\mathfrak{d}^{-1}] &= \frac{1}{\text{Vol}(\mathcal{C}_M)} \int_{c \in \mathcal{C}_M} \frac{1}{\rho_{\varsigma}(x\mathfrak{a} - c)} \underbrace{\sum_{\mathfrak{d} \in \mathcal{I}_K} \sum_{\substack{v \in x\mathfrak{a} \\ (v) = x\mathfrak{a}\mathfrak{d}}} \rho_{\varsigma}(v - c)}_{\rho_{\varsigma}(x\mathfrak{a} - c)} dc \\ &= \frac{1}{\text{Vol}(\mathcal{C}_M)} \int_{c \in \mathcal{C}_M} dc = 1. \end{aligned}$$

Remark 5.4. The instantiation of $\varsigma \in \mathbb{R}_{>0}$ in Definition 5.3 is chosen this way because of the lower bound $\varsigma \geq 2^{n+1}\sqrt{n} \cdot |\Delta_K|^{1/(2n)} \cdot \lambda_n(\mathcal{O}_K)$ and $\lambda_n(\mathcal{O}_K) \geq n\sqrt{|\Delta_K|}$ (see Lemma 2.22). The first of these lower bounds arises

from the size of an LLL-reduced basis of $x\mathfrak{a}$; and the standard deviation ς needs to be larger than this basis size for an efficient computation of the discrete Gaussian over $x\mathfrak{a}$ by Klein's algorithm [GPV08; Kle00].

The instantiation of $M = 2\sqrt{n} \cdot \varsigma$ (or larger) is required in order to have a balanced $v \in K_{\mathbb{R}}$ in line 3 of Algorithm 3. A balanced $v \in K_{\mathbb{R}}$ means that all entries v_{σ} of v are of roughly the same size, i.e., that $\frac{\max_{\sigma} |v_{\sigma}|}{\min_{\sigma} |v_{\sigma}|}$ is small. This has as a consequence that $\mathfrak{d}^{-1} \leftarrow \mathcal{D}_{x\mathfrak{a}}$ and $x\mathfrak{a}$ have a very similar geometry (see Lemma 5.5 part (iii)).

Properties of the distribution representation

Because the distribution $\mathcal{D}_{(x,\mathfrak{a})}$ in Definition 5.3 depends on the ideal lattice $x\mathfrak{a}$ and not on the representing pair (x, \mathfrak{a}) , we can see the domain of the map \mathcal{D} as the group of ideal lattices IdLat_K , i.e., $\mathcal{D} : \text{IdLat}_K \rightarrow L_1(\mathcal{I}_K)$, $x\mathfrak{a} \mapsto \mathcal{D}_{x\mathfrak{a}}$. Even more is true – two *isometric* ideal lattices $x\mathfrak{a}$ and $y\mathfrak{b}$ also have the same distribution $\mathcal{D}_{x\mathfrak{a}}$ and $\mathcal{D}_{y\mathfrak{b}}$. Two ideal lattices being isometric means that there exists an element $\xi = (\xi_{\sigma})_{\sigma} \in \mathcal{C}_1 = \{(y_{\sigma})_{\sigma} \mid |y_{\sigma}| = 1\}$ such that $x\mathfrak{a} = \xi y\mathfrak{b}$ (see Definition 2.20). So, this map can even be interpreted to have domain Pic_K^0 .

Another two remarkable properties of the distribution $\mathcal{D}_{x\mathfrak{a}}$ are that $\mathfrak{d}^{-1} \leftarrow \mathcal{D}_{x\mathfrak{a}}$ always lies in the ideal class $[\mathfrak{a}]$ and has (with high probability) a geometry very similar to $x\mathfrak{a}$. So, in some sense, we may see a sample $\mathfrak{d}^{-1} \leftarrow \mathcal{D}_{x\mathfrak{a}}$ as a sort of ‘discrete approximation’ of the ideal lattice $x\mathfrak{a}$. These important properties of the distribution representation are spelled out in the following lemma.

Lemma 5.5 (Properties of the distribution representation). *The map $\mathcal{D} : \text{IdLat}_K \rightarrow L_1(\mathcal{I}_K)$ has the following properties.*

(i) *(Isometric lattices have the same distribution) For all $x\mathfrak{a}, y\mathfrak{b} \in \text{IdLat}_K$ which are isometric, i.e., $x\mathfrak{a} \sim y\mathfrak{b}$, we have $\mathcal{D}_{x\mathfrak{a}} = \mathcal{D}_{y\mathfrak{b}}$.*

(ii) (*Supported by a single ideal class*) For all $x\mathfrak{a} \in \text{IdLat}_K$, the distribution $\mathcal{D}_{x\mathfrak{a}}$ on \mathcal{I}_K is supported only by inverted integral ideals that lie in the ideal class $[\mathfrak{a}]$.

(iii) (*Bounded size*) For all $x\mathfrak{a} \in \text{IdLat}_K$ with $\mathcal{N}(\mathfrak{a}) \prod_\sigma x_\sigma = 1$, the weight of the distribution $\mathcal{D}_{x\mathfrak{a}}$ is concentrated on inverted integral ideals \mathfrak{d}^{-1} for which holds $\mathcal{N}(\mathfrak{d}^{-1}) \geq (\varsigma + M)^{-n}$. Concretely,

$$\Pr_{\mathfrak{d}^{-1} \leftarrow \mathcal{D}_{x\mathfrak{a}}} [\mathcal{N}(\mathfrak{d}^{-1}) < (\varsigma + M)^{-n}] \leq 2e^{-n}.$$

(iv) (*Similar geometry*) For all $x\mathfrak{a} \in \text{IdLat}_K$ with $\mathcal{N}(\mathfrak{a}) \prod_\sigma x_\sigma = 1$, for almost all $\mathfrak{d}^{-1} \leftarrow \mathcal{D}_{x\mathfrak{a}}$, we have $\mathfrak{d}^{-1} = v^{-1}x\mathfrak{a}$ with $\|v\|_\infty \|v^{-1}\|_\infty \leq 3$, i.e., v is balanced. Concretely,

$$\Pr_{\mathfrak{d}^{-1} \leftarrow \mathcal{D}_{x\mathfrak{a}}} [\exists v \in K_{\mathbb{R}} : \mathfrak{d}^{-1} = v^{-1}x\mathfrak{a} \text{ and } \|v\|_\infty \|v^{-1}\|_\infty \leq 3] \geq 1 - 2e^{-n}$$

Proof. (i) Write $x\mathfrak{a} = \xi y\mathfrak{b}$, use Definition 5.3 and use the fact that $|\xi_\sigma| = 1$ for all embeddings σ to deduce, for a fixed integral ideal \mathfrak{d} ,

$$\begin{aligned} \frac{1}{\rho_\varsigma(x\mathfrak{a} - c)} \sum_{\substack{v \in x\mathfrak{a} \\ (v) = x\mathfrak{a}\mathfrak{d}}} \rho_\varsigma(v - c) &= \frac{1}{\rho_\varsigma(\xi y\mathfrak{b} - c)} \sum_{\substack{v \in y\mathfrak{b} \\ (v) = y\mathfrak{b}\mathfrak{d}}} \rho_\varsigma(\xi v - c) \\ &= \frac{1}{\rho_\varsigma(y\mathfrak{b} - \xi^{-1}c)} \sum_{\substack{v \in y\mathfrak{b} \\ (v) = y\mathfrak{b}\mathfrak{d}}} \rho_\varsigma(v - \xi^{-1}c). \end{aligned}$$

The map $c \mapsto \xi^{-1}c$ is an isometric smooth bijection on the hypercircle \mathcal{C}_M , so integrating with respect to the variable $\xi^{-1}c$ or c for $c \in \mathcal{C}_M$ doesn't change the value of the integral. Therefore, $\mathcal{D}_{(x,\mathfrak{a})}[\mathfrak{d}^{-1}] = \mathcal{D}_{(y,\mathfrak{b})}[\mathfrak{d}^{-1}]$ for all $\mathfrak{d} \in \mathcal{I}_K$.

(ii) From Equation (5.65) we see that $\mathfrak{d} = (v)/(x\mathfrak{a})$ for $v \in x\mathfrak{a}$ and therefore \mathfrak{d} must be an integral ideal in the inverse class of \mathfrak{a} ; so \mathfrak{d}^{-1} lies in the ideal class $[\mathfrak{a}]$.

(iii) Use the fact that $\mathcal{N}(x\mathfrak{a}) = 1$ and the fact that $(v) = x\mathfrak{a}\mathfrak{d}$ to derive

$$\begin{aligned} \Pr_{\mathfrak{d} \leftarrow \mathcal{D}_{x\mathfrak{a}}} [\mathcal{N}(\mathfrak{d}) > (\varsigma + M)^n] &= \Pr_{\mathfrak{d} \leftarrow \mathcal{D}_{x\mathfrak{a}}} [\mathcal{N}(x\mathfrak{a}\mathfrak{d}) > (\varsigma + M)^n] \\ &\leq \Pr_{\substack{c \leftarrow \mathcal{C}_M \\ v \leftarrow \mathcal{G}_{x\mathfrak{a}, \varsigma, c}}} [\|v\|_2 > \sqrt{n} \cdot \varsigma + \sqrt{n} \cdot M]. \\ &\leq \max_{c \in \mathcal{C}_M} \Pr_{v \leftarrow \mathcal{G}_{x\mathfrak{a}, \varsigma, c}} [\|v - c\|_2 > \sqrt{n} \cdot \varsigma]. \end{aligned}$$

Where the first inequality follows from norm inequalities; we have $n^{-1/2} \cdot \|v\|_2 \geq n^{-1} \cdot \|v\|_1 \geq \mathcal{N}(v)^{1/n} = \mathcal{N}(x\mathfrak{a}\mathfrak{d})^{1/n} \geq (\varsigma + M)$. The second inequality follows from the triangle inequality and the fact that $\|c\| = \sqrt{n} \cdot M$. By Banaszczyk's tail bound (see Lemma 2.25) and by smoothing arguments (see Lemma 2.31), we conclude

$$\Pr_{v \leftarrow \mathcal{G}_{x\mathfrak{a}, \varsigma, c}} [\|v - c\| \geq \sqrt{n} \cdot \varsigma] \leq e^{-n} \cdot \frac{\rho_\varsigma(x\mathfrak{a})}{\rho_\varsigma(x\mathfrak{a} - c)} \leq 2e^{-n}.$$

For the smoothing argument we use the fact that $\varsigma = 2^{n+1} \cdot n \cdot |\Delta_K| \geq n\lambda_n(\mathcal{O}_K) \geq \eta_1(x\mathfrak{a})$ (see [MR07, Lm. 3.3 and 3.4]).

(iv) We have, by the norm inequalities, $\|v - c\| \geq \|v - c\|_\infty$, and therefore, by part (iii) of this lemma,

$$\Pr_{v \leftarrow \mathcal{G}_{x\mathfrak{a}, \varsigma, c}} [\|v - c\|_\infty \geq \sqrt{n} \cdot \varsigma] \leq \Pr_{v \leftarrow \mathcal{G}_{x\mathfrak{a}, \varsigma, c}} [\|v - c\| \geq \sqrt{n} \cdot \varsigma] \leq 2e^{-n}.$$

Since $|c_\sigma| = M$ for all embeddings σ , and since $M = 2\sqrt{n}\varsigma$, we have, except with probability $2e^{-n}$,

$$\|v\|_\infty \|v^{-1}\|_\infty = \frac{\max_\sigma |v_\sigma|}{\min_\sigma |v_\sigma|} \leq \frac{M + \sqrt{n}\varsigma}{M - \sqrt{n}\varsigma} = 3.$$

□

Remark 5.6. The bound $\|v\|_\infty \|v^{-1}\| \leq 3$ w.h.p. in part (iv) of Lemma 5.5 can be tightened to $\|v\|_\infty \|v^{-1}\| \leq 1 + O(e^{-n})$ by taking $M = 2^n \cdot \varsigma$. Because this would only remove a rather non-significant constant 3 in the quality loss of the output in the worst-case to average-case reduction, we choose this 'constant bound' for simplicity.

A consequence of Lemma 5.5 is that the map $\mathcal{D} : \text{IdLat}_K \rightarrow L_1(\mathcal{I}_K)$, that sends ideal lattices to distributions on \mathcal{I}_K factors through the quotient group IdLat_K / \sim , where ‘ \sim ’ stands for factoring out by isometries. As the group of ideal lattices up to isometries is naturally isomorphic to the Arakelov class group Pic_K^0 (see Lemma 2.21), we might as well consider Pic_K^0 as the domain of the map \mathcal{D} .

5.4. The Worst-case to Average-case Reduction

Introduction

A worst-case to average-case reduction consists of two main parts: the definition of the average-case distribution and an algorithm that reduces a fixed problem instance to a sample of the average-case distribution.

We start this section with the definition of the average-case distribution, which is derived from the uniform distribution on the Arakelov class group. After that, we will describe the reduction algorithm. In this description of the worst-case to average-case reduction we temporarily ignore issues regarding real numbers. In the last part of this section we will prove the correctness of the reduction algorithm and examine the precise quality loss that occurs in the reduction.

Discussing and solving the issues regarding real numbers and finite precision in the distribution algorithm Algorithm 3 and the reduction algorithm Algorithm 4 is deferred to Section 5.5. In that section we will prove that both the distribution algorithm and the reduction algorithm can be run efficiently on a finite machine by means of appropriate discretization.

Definition of the average-case distribution

Knowing in advance that the reduction algorithm will make use of the random walk machinery of Chapter 4, which leads to a near-uniform distribution

on the Arakelov class group, the average-case distribution must be strongly tied to this distribution.

Indeed, we define the average-case distribution to be distribution on $L_1(\mathcal{I}_K)$ defined by the following rule.

$$\mathcal{D}_{\mathcal{U}(\text{Pic}_K^0)}[\mathfrak{d}^{-1}] = \frac{1}{|\text{Pic}_K^0|} \int_{\mathbf{a} \in \text{Pic}_K^0} \mathcal{D}_{\mathbf{a}}[\mathfrak{d}^{-1}] d\mathbf{a}. \quad (5.66)$$

In essence this is just ‘taking the average’ of all distributions $\mathcal{D}_{\mathbf{a}}$ (as in Section 5.3) where \mathbf{a} is taken uniformly from the Arakelov class group.

Reduction algorithm

The reduction algorithm essentially consists of taking an input ideal lattice $x\mathfrak{a}$, applying a specific random walk procedure on it as in Chapter 4, yielding $\tilde{x}\tilde{\mathfrak{a}}$, and sampling an ideal $\mathfrak{d}^{-1} \leftarrow \mathcal{D}_{\tilde{x}\tilde{\mathfrak{a}}}$. A rigorous, precise description of this procedure is spelled out in Algorithm 4.

Remark 5.7. *Algorithm 4, and also Algorithm 3, are strictly spoken not algorithms that can be run on a finite computer, because of the continuous distributions occurring in the algorithm descriptions. In Algorithm 3 it is the uniform sampling from the hypercircle \mathcal{C}_M and in Algorithm 4 it is the Gaussian sampling that is inherently continuous.*

In Section 5.5 we will show that those continuous distributions can be efficiently discretized without a significant impact on the final result. Therefore, we just ignore these continuity issues for now, for the sake of clarity and brevity.

Explanation of the reduction algorithm

Algorithm 4: The worst-case to average-case reduction algorithm

Require:

- A pair $(x, \mathfrak{a}) \in K_{\mathbb{R}} \times \mathcal{I}_K$ satisfying $\mathcal{N}(\mathfrak{a}) \prod_{\sigma} x_{\sigma} = 1$.
- The values $[\Lambda_K : C]$ and $\eta_1(C^*)$ of a suitable sublattice $C \subseteq \Lambda_K$ of the logarithmic unit lattice,
- An oracle \mathcal{A} that solves γ -Hermite SVP in \mathfrak{d}^{-1} whenever $\mathfrak{d}^{-1} \leftarrow \mathcal{D}_{\mathcal{U}(\text{Pic}_K^0)}$.

Ensure: A vector $\alpha \in x\mathfrak{a}$ that is a solution to $B^{1/n} \cdot \gamma$ -Hermite SVP in the ideal lattice $x\mathfrak{a}$, i.e.,

$$\|\alpha\| \leq \gamma \cdot B^{1/n} \cdot \det(x\mathfrak{a})^{1/n},$$

where $B = \tilde{O}(4^n \cdot \eta_1(C^*)^{\mathfrak{r}} \cdot [\Lambda_K : C] \cdot |\text{Pic}_K^0| \cdot (\log |\Delta_K|)^2)$, or, *failure*.

- 1: Put $s = \max(\sqrt{2} \cdot \eta_1(C^*), (\log n)^2)$ and $B = \tilde{O}(4^n \cdot \eta_1(C^*)^{\mathfrak{r}} \cdot [\Lambda_K : C] \cdot |\text{Pic}_K^0| \cdot (\log |\Delta_K|)^2)$ as in Corollary 5.8.
- 2: Multiply the ideal \mathfrak{a} by a prime ideal \mathfrak{p} uniformly sampled from the set $\{\mathfrak{p} \text{ prime ideal of } \mathcal{O}_K \mid \mathcal{N}(\mathfrak{p}) \leq B\}$, yielding $\mathfrak{a}\mathfrak{p}$.
- 3: Sample a Gaussian distributed $y \leftarrow \mathcal{G}_{s, H}$, where H is the hyperplane where the logarithmic unit lattice lives in.
- 4: Put $p = \mathcal{N}(\mathfrak{p})^{1/n}$, so that $e^y x \mathfrak{a} \mathfrak{p} / p$ has norm 1, where $e^y \in K_{\mathbb{R}}$ is the component-wise exponentiation of $y \in H$.
- 5: Sample $\mathfrak{d}^{-1} \leftarrow \mathcal{D}_{e^y \cdot x \cdot \mathfrak{a} \mathfrak{p} / p}$ using Algorithm 3, and let $v \in e^y x \mathfrak{a} \mathfrak{p} / p$ be the additional output of Algorithm 3 that satisfies $\mathfrak{d}^{-1} = v^{-1} e^y x \mathfrak{a} \mathfrak{p} / p$.
- 6: Invoke the γ -Hermite SVP oracle \mathcal{A} on $\mathcal{D}_{\mathcal{U}(\text{Pic}_K^0)}$ to find a $\kappa \in \mathfrak{d}^{-1}$ for which holds $\|\kappa\| \leq \gamma \cdot \det(\mathfrak{d}^{-1})^{1/n}$
- 7: **return** $p \cdot e^{-y} \cdot v \cdot \kappa \in x\mathfrak{a}$.

Randomize the input ideal lattice $x\mathfrak{a}$. The first four steps of Algorithm 4 actually applies a random walk on the input ideal lattice $x\mathfrak{a}$, resulting in a randomized ideal lattice $\tilde{x}\tilde{\mathfrak{a}}$. By the results of Chapter 4, this ideal lattice is nearly uniformly distributed in the Arakelov class group. Therefore, sampling from the distribution $\mathcal{D}_{\tilde{x}\tilde{\mathfrak{a}}} \approx \mathcal{D}_{\mathcal{U}(\text{Pic}_K^0)}$ associated with this random ideal lattice $\tilde{x}\tilde{\mathfrak{a}}$ then yields an ideal \mathfrak{d}^{-1} that must be closely distributed as $\mathcal{D}_{\mathcal{U}(\text{Pic}_K^0)}$. So, that is an intuitive idea of why the output ideal \mathfrak{d}^{-1} is almost distributed as the *average-case distribution* as in Equation (5.66).

Sample $\mathfrak{d}^{-1} \leftarrow \mathcal{D}_{\tilde{x}\tilde{\mathfrak{a}}}$ and apply the Hermite-SVP oracle \mathcal{A} on \mathfrak{d}^{-1} . Because $\mathfrak{d}^{-1} \leftarrow \mathcal{D}_{\tilde{x}\tilde{\mathfrak{a}}}$ is so close to the average-case distribution, we can actually invoke the oracle \mathcal{A} to find a short vector in the ideal $\mathfrak{d}^{-1} \leftarrow \mathcal{D}_{\tilde{x}\tilde{\mathfrak{a}}}$. The sampling is done in step 5 and calling the oracle in step 6 of Algorithm 4.

Transform the short vector $\gamma \in \mathfrak{d}^{-1}$ into a short vector in the randomized ideal lattice $\tilde{x}\tilde{\mathfrak{a}}$. Recall that Algorithm 3 on input $\tilde{x}\tilde{\mathfrak{a}}$ outputs both $\mathfrak{d}^{-1} \leftarrow \mathcal{D}_{\tilde{x}\tilde{\mathfrak{a}}}$ and a $v \in \tilde{x}\tilde{\mathfrak{a}}$ such that $\mathfrak{d}^{-1}v = \tilde{x}\tilde{\mathfrak{a}}$.

So, any short vector $\kappa \in \mathfrak{d}^{-1}$ can be transformed into a short vector $v\kappa \in \tilde{x}\tilde{\mathfrak{a}}$. Because this $v \in \tilde{x}\tilde{\mathfrak{a}}$ is *balanced*, this does not affect the shortness of the vector κ much; in a way one might say that the ideal lattices $\tilde{x}\tilde{\mathfrak{a}}$ and \mathfrak{d}^{-1} geometrically very much resemble each other.

Transform a short vector in $\tilde{x}\tilde{\mathfrak{a}}$ to a short vector in the input ideal lattice $x\mathfrak{a}$. By construction, $\tilde{x}\tilde{\mathfrak{a}} = e^y/p \cdot x\mathfrak{a}\mathfrak{p}$, i.e., the randomized ideal lattice is just the input ideal lattice multiplied by a prime ideal, slightly disturbed a renormalized. By undoing the disturbance (i.e., dividing by e^y) and undoing the renormalization (i.e., multiplying by $p = \mathcal{N}(\mathfrak{p})^{1/n}$) on the short vector $v\kappa \in \tilde{x}\tilde{\mathfrak{a}}$, we obtain a short vector in $x\mathfrak{a}\mathfrak{p} \subseteq x\mathfrak{a}$. More precisely: because $v\kappa \in \tilde{x}\tilde{\mathfrak{a}}$, we have that $p \cdot e^{-y} \cdot (v\kappa) \in x\mathfrak{a}\mathfrak{p} \subseteq x\mathfrak{a}$.

Reason for quality loss. Note that the reduction algorithm only ensures to find a vector solving $B^{1/n} \cdot \gamma$ -Hermite SVP, whereas the oracle \mathcal{A} in

Algorithm 4 is assumed to be able to find a vector satisfying γ -Hermite SVP on an ‘average case’ ideal lattice (see Equation (5.66)).

This particular loss $B^{1/n}$ comes from the fact that we cannot not reasonably ‘undo’ the part of the random walk where we multiply the input ideal lattice $x\mathfrak{a}$ by a random prime ideal \mathfrak{p} . So, this reduction algorithm actually finds a γ -Hermite short vector in $x\mathfrak{a}\mathfrak{p}$, a slightly wider ideal than $x\mathfrak{a}$. As $x\mathfrak{a}\mathfrak{p} \subseteq x\mathfrak{a}$, and the root determinants of these ideal lattices differ with a factor $p = \mathcal{N}(\mathfrak{p})^{1/n}$, a γ -Hermite short vector in $x\mathfrak{a}\mathfrak{p}$ is a $B^{1/n} \cdot \gamma$ -Hermite short vector in $x\mathfrak{a}$, as $\mathcal{N}(\mathfrak{p})^{1/n} \leq B^{1/n}$.

Proof of correctness and quantification of the quality loss

In order to prove the result of this chapter, we need the following specialization of the random walk theorem of Chapter 4, which is specifically tailored to the worst-case to average-case reduction.

Corollary 5.8 (Random walk in the Arakelov class group, simplified). *Let K be a number field, and let $C \subseteq \Lambda_K$ be a sublattice of the logarithmic unit lattice. Assuming the Extended Riemann Hypothesis, there exists a bound $B = \tilde{O}(4^n \cdot \eta_1(C^*)^r \cdot [\Lambda_K : C] \cdot |\text{Pic}_K^0| \cdot (\log |\Delta_K|)^2)$ such that the random walk distribution with one step $\mathcal{W}_{\text{Pic}_K^0}(B, 1, s)$ is exponentially close to uniform in $L_1(\text{Pic}_K^0)$.*

$$\|\mathcal{W}_{\text{Pic}_K^0}(B, 1, s) - \mathcal{U}(\text{Pic}_K^0)\|_1 \leq 2^{-n}$$

Proof. Apply Theorem 4.18 from Chapter 4 with

- $k = \frac{1}{2\log n} \cdot (r \cdot \log(1/\tilde{s}) + \log(\text{Vol}(\text{Pic}_K^0)) + 2\log(1/\varepsilon) + \log[\Lambda_K : C] + 2)$, so that taking $N = 1$ satisfies the requirements of the theorem.
- $s = 1/(\sqrt{2} \cdot \eta_1(C^*))$, so that $\tilde{s} = 1/\eta_1(C^*)$;
- $\varepsilon = 2^{-n}$.

Appropriately substituting above instantiations in the formula for B in Theorem 4.18, noting that $n^{2k} = O(\eta_1(C^*)^r \cdot |\text{Pic}_K^0| \cdot 4^n \cdot [\Lambda_K : C])$, we obtain

$$\begin{aligned} B &= \tilde{O}\left(n^{2k}[n^2(\log \log(1/\varepsilon))^2 + n^2(\log(1/\tilde{s}))^2 \right. \\ &\quad \left. + n^2 \log([\Lambda_K : C])^2 + (\log |\Delta_K|)^2]\right) \\ &= \tilde{O}(4^n \cdot \eta_1(C^*)^r \cdot [\Lambda_K : C] \cdot |\text{Pic}_K^0| \cdot (\log |\Delta_K|)^2). \end{aligned}$$

□

Using above specialized random walk theorem, we can prove the main theorem of this chapter.

Theorem 5.9. *Let K be a number field with logarithmic unit lattice Λ_K , let $C \subseteq \Lambda_K$ be any sublattice, and denote its dual lattice by C^* . Put $1/s = \max(\sqrt{2} \cdot \eta_1(C^*), \log(n)^2)$.*

Assume we have a (possibly randomized) algorithm \mathcal{A} that solves γ -Hermite-SVP within an approximation factor $\gamma \geq 1$ and probability¹ at least $q > 0$ when given an input \mathfrak{a} with $\mathfrak{a} \leftarrow \mathcal{D}_{\mathcal{U}(\text{Pic}_K^0)}$.

Then there exists a randomized algorithm \mathcal{B} solving $(O(B^{1/n}) \cdot \gamma)$ -Hermite-SVP in any ideal lattice $x\mathfrak{a} \in \text{IdLat}_K$, with probability² at least $q - n^{-\omega(1)}$, where $B = \tilde{O}(4^n \cdot s^{-r} \cdot [\Lambda_K : C] \cdot |\text{Pic}_K^0| \cdot (\log |\Delta_K|)^2)$. The algorithm \mathcal{B} runs within time polynomial in $\log |\Delta_K|, \log [\Lambda_K : C], \text{size}(x)$ and $\text{size}(M_{\mathfrak{a}})$ on input $(x, \mathfrak{a}) \in K_{\mathbb{R}} \times \mathcal{I}_K$ and needs one call to the algorithm \mathcal{A} .

¹Here, the probability q is taken over the random choice of $\mathfrak{a} \leftarrow \mathcal{D}_{\mathcal{U}(\text{Pic}_K^0)}$ and over the possible internal randomness of the algorithm \mathcal{A}

²Here, the probability is taken over the internal randomness of \mathcal{B}

Furthermore, the loss $B^{1/n}$ in the approximation factor of Hermite-SVP in the reduction can be upper bounded as follows.

$$B^{1/n} = \begin{cases} \tilde{O}(\sqrt{n}) & \text{if } K = \mathbb{Q}(\zeta_{p^k}), \text{ a prime power cyclotomic field, assuming that } h_K^+ = \log(n)^{O(n)}. \\ \tilde{O}(n^{1-n_{\mathbb{C}}/n} \cdot |\Delta_K|^{1/(2n)}) & \text{otherwise} \end{cases} \quad (5.67)$$

Proof. Proof of the running time. The random walk process and the distribution representation of the reduction have inherently continuous aspects, that need to be discretized in order to be suitable for an actual computer. The discretized version of the reduction is treated in Section 5.5, in which also its running time and its discretization error is studied. In Theorem 5.11 we show that the reduction can be approximated within a negligible error margin, using time polynomial in $\log |\Delta_K|, \log[\Lambda_K : C], \text{size}(x)$; here we take $\varepsilon = 2^{-n}$ to have exponentially small error.

Success probability. By the choice of parameters in reduction Algorithm 4, the Arakelov class of $xe^y/p \cdot \mathfrak{ap}$ (where \mathfrak{p} and $y \in H$ are randomly chosen as in Algorithm 4) must be exponentially close to uniform in Pic_K^0 in total variation distance (see Corollary 5.8). By the data processing inequality [CT06, §2.8], $\mathcal{D}_{xe^y/p \cdot \mathfrak{ap}}$ is exponentially close to $\mathcal{D}_{\mathcal{U}(\text{Pic}_K^0)}$ as well. Therefore, the algorithm \mathcal{A} cannot distinguish reasonably between the two distributions and outputs with probability at least $q - 2^{-n}$ a solution of γ -Hermite SVP in $\mathfrak{d}^{-1} \leftarrow \mathcal{D}_{xe^y/p \cdot \mathfrak{ap}}$.

Quality of the output. Let us assume that algorithm \mathcal{A} indeed found a solution to γ -Hermite SVP, i.e., a vector $\kappa \in \mathfrak{d}^{-1}$ which satisfies $\|\kappa\| \leq \gamma \cdot \det(\mathfrak{d}^{-1})^{1/n}$, where $\mathfrak{d}^{-1} \leftarrow \mathcal{D}_{xe^y/p \cdot \mathfrak{ap}}$.

As $\kappa \in \mathfrak{d}^{-1} = v^{-1}e^y x/p \cdot \mathfrak{ap}$ (see Algorithm 3), we must have that³ $\kappa = v^{-1}e^y/p \cdot \alpha$ for some $\alpha \in x\mathfrak{ap}$. This particular $\alpha \in x\mathfrak{ap} \subset x\mathfrak{a}$ is a solution for $O(B^{1/n}) \cdot \gamma$ -Hermite SVP in $x\mathfrak{a}$, which can be seen by the following

³Note that $p = \mathcal{N}(\mathfrak{p})^{1/n}$

lines of reasoning. We have the following bound on $\|\alpha\|$, where we write out $p = \mathcal{N}(\mathfrak{p})^{1/n}$,

$$\begin{aligned}\|\alpha\| &= \|ve^{-y}\kappa\| \cdot \mathcal{N}(\mathfrak{p})^{1/n} \leq \|v\|_\infty \|e^{-y}\|_\infty \|\kappa\| \cdot \mathcal{N}(\mathfrak{p})^{1/n} \\ &\leq \|v\|_\infty \|e^{-y}\|_\infty \cdot \gamma \cdot \det(\mathfrak{d}^{-1})^{1/n} \cdot \mathcal{N}(\mathfrak{p})^{1/n},\end{aligned}\tag{5.68}$$

But also, by the fact that multiplication by e^y doesn't change the determinant and $\det(x/p \cdot \mathfrak{a}\mathfrak{p}) = \det(x\mathfrak{a})$ (by definition of $p = \mathcal{N}(\mathfrak{p})^{1/n}$), we have

$$\det(\mathfrak{d}^{-1}) = \det(v^{-1}e^y x/p \cdot \mathfrak{a}\mathfrak{p}) \leq \|v^{-1}\|_\infty^n \cdot \det(x\mathfrak{a}).\tag{5.69}$$

Combining Equation (5.68) and Equation (5.69), using the fact that $\mathcal{N}(\mathfrak{p}) \leq B$, and $\|v^{-1}\|_\infty \|v\|_\infty \leq 3$ with high probability (see Lemma 5.5, ‘ v is balanced’), we obtain

$$\begin{aligned}\|\alpha\| &\leq \underbrace{\|v\|_\infty \cdot \|v^{-1}\|_\infty}_{\leq 3 \text{ (w.h.p.)}} \cdot \underbrace{\|e^{-y}\|_\infty}_{\leq 3 \text{ (w.h.p.)}} \cdot \underbrace{\mathcal{N}(\mathfrak{p})^{1/n}}_{\leq B^{1/n}} \cdot \gamma \cdot \det(x\mathfrak{a})^{1/n} \\ &\leq 9 \cdot B^{1/n} \cdot \gamma \cdot \det(x\mathfrak{a})^{1/n}.\end{aligned}$$

Here, the bound on $\|e^y\|_\infty$ can be obtained by the fact that $y \leftarrow \mathcal{G}_{H,s}$ is from a Gaussian distribution, with⁴ $s \leq 1/(\log n)^2$. Namely, $\|y\|_\infty \leq (\log n)^2 s \leq 1$ except with probability at most $2^{-\Omega((\log n)^2)} = n^{-\omega(1)}$. Therefore $\|e^y\|_\infty \leq e^{\|y\|_\infty} \leq 3$ except with probability $n^{-\omega(1)}$.

Conclusion. So, with probability $q - n^{-\omega(1)}$, algorithm \mathcal{B} solves $9 \cdot B^{1/n} \cdot \gamma$ -Hermite SVP in the input ideal lattice $x\mathfrak{a} \in \text{IdLat}_K$, within polynomial time in $\log |\Delta_K|, \log [\Lambda_K : C]$, $\text{size}(x)$ and $\text{size}(M_{\mathfrak{a}})$, and using one call to the algorithm \mathcal{A} . The explicit bounds on $B^{1/n}$ in Equation (5.67) are proved in Proposition 5.10. \square

⁴Note that $s \leq 1/(\log n)^2$, by the instantiation $1/s = \max(\sqrt{2} \cdot \eta_1(C^*), \log(n)^2)$ in the theorem.

Proposition 5.10. *The loss $B^{1/n}$ in the approximation factor of Hermite-SVP in the reduction of Theorem 5.9 can be upper bounded as follows.*

$$B^{1/n} = \begin{cases} \tilde{O}(\sqrt{n}) & \text{if } K = \mathbb{Q}(\zeta_{p^k}), \text{ a prime power cyclotomic field, assuming that } h_K^+ = \log(n)^{O(n)}. \\ \tilde{O}(n^{1-n_{\mathbb{C}}/n} \cdot |\Delta_K|^{1/(2n)}) & \text{otherwise} \end{cases}$$

Proof. The difference between the upper bounds of $B^{1/n}$ for different types of number fields depends on the choice of the sublattice $C \subseteq \Lambda_K$ of the logarithmic unit lattice. Because $1/s = \max(\sqrt{2}\eta_1(C^*), \log(n)^2)$, the product $s^{-\mathbb{r}} \cdot [\Lambda_K : C]$ is the only part of B that depends on the choice of this sublattice.

For general number fields, we will choose $C = \Lambda_K$, and use a general upper bound $\eta_1(\Lambda_K^*) \leq O(n(\log n)^3)$ due to Kessler and Dobrowolski [Kes91; Dob79] to obtain $s^{-\mathbb{r}} \cdot [\Lambda_K : C] \leq O(n^{\mathbb{r}} \log(n)^{3\mathbb{r}})$.

For cyclotomic number fields with prime power conductor, we choose $C \subseteq \Lambda_K$ to be the sublattice of Λ_K consisting of the logarithmic image of the cyclotomic units [Was12, Ch. 8]. For this sublattice it is known that $[\Lambda_K : C] = h_K^+$, the class number of the maximal totally real subfield of K , and $\eta_1(C^*) \leq O(1)$, so that $s^{-\mathbb{r}} \cdot [\Lambda_K : C] \leq O(\log(n)^{2\mathbb{r}} \cdot h_K^+) = \log(n)^{O(n)}$ for these prime power cyclotomic number fields, under the assumption that $h_K^+ = \log(n)^{O(n)}$. The precise derivation of these bounds follow later in this proof.

Plugging these bounds into the value of B in Theorem 5.9, using $\mathbb{r} = n - n_{\mathbb{C}} - 1 \leq n$, $|\text{Pic}_K^0|^{1/n} = \tilde{O}(|\Delta_K|^{1/(2n)})$ (see Lemma 2.17), $|\Delta_K|^{1/(2n)} \leq \sqrt{n}$

for cyclotomic fields, and suppressing polylogarithmic factors, we obtain

$$\begin{aligned}
B^{1/n} &= \tilde{O}\left(\underbrace{s^{-\mathfrak{r}/n} \cdot [\Lambda_K : C]^{1/n}}_{=\log(n)^{O(1)} \text{ for prime power cyclotomic fields}} \cdot \underbrace{\tilde{O}(|\Delta_K|^{1/(2n)})}_{\text{for general number fields}} \cdot \underbrace{(\log |\Delta_K|)^{2/n}}_{\text{polylog. factor}}\right) \\
&= \begin{cases} \tilde{O}(\sqrt{n}) & \text{for prime power cyclotomic fields,} \\ & \text{assuming that } h_K^+ = \log(n)^{O(n)} \\ \tilde{O}(n^{1-n_{\mathbb{C}}/n} \cdot |\Delta_K|^{1/(2n)}) & \text{for general number fields} \end{cases}
\end{aligned}$$

General number fields. We take $C = \Lambda_K$, so that $[\Lambda_K : C] = 1$. By the fact that $\eta_1(\Lambda_K^*) \leq \frac{\sqrt{\mathfrak{r}}}{\lambda_1(\Lambda_K)}$ [MR07, Lm. 3.2] and by the general upper bound $1/\lambda_1(\Lambda_K) \leq 1000\sqrt{\mathfrak{r}+1} \log(\mathfrak{r})^3$ [Kes91; Dob79], we obtain $\eta_1(\Lambda_K^*) \leq \sqrt{\mathfrak{r}}/\lambda_1(\Lambda_K) \leq 2000 \cdot \mathfrak{r} \cdot \log(\mathfrak{r})^3$. Therefore, since $1/s = \max(\sqrt{2} \cdot \eta_1(C^*), (\log n)^2)$,

$$s^{-\mathfrak{r}} \cdot [\Lambda_K : C] \leq O(n^{\mathfrak{r}} \log(n)^{3\mathfrak{r}}) \text{ for general number fields } K$$

Prime power cyclotomic number fields. We take C to be the logarithmic image of the *group of cyclotomic units*, which are units that have a specific compact shape [Was12, Ch. 8]. For this *logarithmic cyclotomic unit lattice* $C \subseteq \Lambda_K$, holds $[\Lambda_K : C] = h_K^+$, the class number of the maximal real field in the cyclotomic field K [Was12, Thm. 8.2]. Due to a result of Cramer et al. [Cra+16, Thm. 3.1] we have an upper bound on the last successive minimum $\lambda_{\mathfrak{r}}(C^*)$ of the dual logarithmic cyclotomic unit lattice. Combined with a general smoothing parameter bound for lattices [MR07, Lm. 3.3], this yields the following bound on the smoothing parameter of the dual logarithmic cyclotomic unit lattice: $\eta_1(C^*) \leq \log(4\mathfrak{r})\lambda_{\mathfrak{r}}(C^*) \leq O(\log(\mathfrak{r})^{5/2} \cdot \mathfrak{r}^{-1/2}) = O(1)$. Therefore, with the instantiation $1/s = \max(\sqrt{2} \cdot \eta_1(C^*), (\log n)^2)$,

$$s^{-\mathfrak{r}} \cdot [\Lambda_K : C] \leq O(\log(n)^{2\mathfrak{r}} \cdot h_K^+) \text{ for prime power cyclotomic fields } K.$$

□

5.5. Discretizing the Reduction Algorithm

5.5.1. Introduction

In the reduction algorithm of Section 5.4 (see Algorithm 4), we saw that the random walk procedure is inherently continuous, due to its continuous Gaussian walk. On top of that, the computation of the distribution representation \mathcal{D} also has a continuous aspect, namely the sampling of a vector on a large circle \mathcal{C}_M .

The purpose of this section is to show that the result of applying the random walk procedure and the distribution representation using only *finite precision* doesn't differ too much from the result when one would use infinite precision instead. In other words, actually computing the random walk and the distribution on a finite machine (as in Algorithm 6 and Algorithm 5) doesn't spoil the end result. In particular, none of the operations in this section involves real numbers; it is all floating point arithmetic.

Additionally, this section also provides an upper bound on the running time of this discretized reduction algorithm.

We define $\mathcal{D}_{\mathcal{W}(\text{Pic}_K^0) + x\mathfrak{a}}$ by the distribution of \mathfrak{d}^{-1} in step 5 of Algorithm 4, and $\ddot{\mathcal{D}}_{\ddot{\mathcal{W}}(\text{Pic}_K^0) + x\mathfrak{a}}$ by the distribution of \mathfrak{d}^{-1} in step 6 of Algorithm 6. A precise description of these distributions for the case $x\mathfrak{a} = \mathcal{O}_K$ can be found in Definition 5.14 and Definition 5.16, respectively. These distributions are only being described for the case $x\mathfrak{a} = \mathcal{O}_K$, as the general case is a mere translation of this base case. Note the dots above $\ddot{\mathcal{D}}$ and $\ddot{\mathcal{W}}$ to indicate discreteness.

Theorem 5.11. *Let $x\mathfrak{a} \in \text{IdLat}_K^0$ be a norm-one ideal lattice, where \mathfrak{a} is represented by a finite-precision matrix $M_{\mathfrak{a}}$ and $x \in K_{\mathbb{R}}$ is represented by a finite-precision vector. Then, Algorithm 6 approximates the distribution of Algorithm 4 within a total variation distance of $23 \cdot \varepsilon$, i.e.,*

$$\|\mathcal{D}_{\mathcal{W}(\text{Pic}_K^0) + x\mathfrak{a}} - \ddot{\mathcal{D}}_{\ddot{\mathcal{W}}(\text{Pic}_K^0) + x\mathfrak{a}}\| \leq 23 \cdot \varepsilon,$$

and runs within time polynomial in $\log |\Delta_K|, \text{size}(x), \text{size}(M_{\mathfrak{a}})$ (see Section 2.1) and $\log(1/\varepsilon)$.

Roadmap of the proof

Introduction. In this proof, we show that the the random walk distribution $\mathcal{D}_{\mathcal{W}(\text{Pic}_K^0) + x\mathfrak{a}}$ from \mathfrak{d}^{-1} in line 5 of Algorithm 4 and the *discretized* random walk distribution $\ddot{\mathcal{D}}_{\ddot{\mathcal{W}}(\text{Pic}_K^0) + x\mathfrak{a}}$ from \mathfrak{d}^{-1} in line 5 of Algorithm 6 are close to each other in the total variation distance.

In the proof we will, without loss of generality, assume that $x\mathfrak{a} = \mathcal{O}_K$. The case of general $x\mathfrak{a}$ consists of a mere translation of the distributions involved and does not affect the proof structure. Therefore, we resort to proving closeness of $\mathcal{D}_{\mathcal{W}(\text{Pic}_K^0)}$ and $\ddot{\mathcal{D}}_{\ddot{\mathcal{W}}(\text{Pic}_K^0)}$.

The proof of closeness in total variation distance proceeds by two steps; the first step discretizes the continuous Gaussian sampling in the reduction Algorithm 4, whereas the second step discretizes the uniform sampling on the M -circle in the distribution Algorithm 3 which is used in the reduction.

Sampling the Gaussian walk in a discrete manner doesn't spoil the resulting distribution. In the random walk procedure, a Gaussian distribution is sampled in the logarithmic unit lattice ambient vector space and subsequently exponentiated component-wise to act on the processed input ideal lattice. This part is referred to as the ‘continuous walk’ of the random walk procedure. A finite computer cannot sample from continuous distributions, so in the actual algorithmic implementation a *discrete Gaussian* is sampled on a sufficiently fine grid – a lattice – on the ambient vector space.

The discrete random walk distribution resulting from sampling the Gaussian walk in this discrete way, whereas keeping the rest of the random walk procedure the same, is what we will call $\ddot{\mathcal{W}}(\text{Pic}_K^0)$. By a technical computation, we will show that $\mathcal{D}_{\mathcal{W}(\text{Pic}_K^0)} \approx \mathcal{D}_{\ddot{\mathcal{W}}(\text{Pic}_K^0)}$.

Sampling from a discrete circle doesn't change the map \mathcal{D} much. In the beginning of the distribution representation \mathcal{D} , a vector is uniformly sampled from a M -circle \mathcal{C}_M in $K_{\mathbb{R}}$. In reality, on a finite computer, we need to sample from this M -circle in a discrete manner, while keeping the rest of the distribution computation the same. This particular map is called $\ddot{\mathcal{D}} : \text{Pic}_K^0 \rightarrow L_1(\mathcal{I}_K)$.

By showing that $\ddot{\mathcal{D}}$ and \mathcal{D} are close for any $\mathbf{a} \in \text{Pic}_K^0$, we draw the conclusion that for any distribution \mathcal{P} on Pic_K^0 , $\ddot{\mathcal{D}}_{\mathcal{P}}$ and $\mathcal{D}_{\mathcal{P}}$ are close as well. In particular, $\ddot{\mathcal{D}}_{\mathcal{W}(\text{Pic}_K^0)} \approx \mathcal{D}_{\mathcal{W}(\text{Pic}_K^0)}$.

Finalizing. By using the above two parts, we can show that the following three distributions are actually close.

$$\mathcal{D}_{\mathcal{W}(\text{Pic}_K^0)} \underset{\substack{\approx \\ \text{First part}}}{\sim} \mathcal{D}_{\ddot{\mathcal{W}}(\text{Pic}_K^0)} \underset{\substack{\approx \\ \text{Second part}}}{\sim} \ddot{\mathcal{D}}_{\ddot{\mathcal{W}}(\text{Pic}_K^0)}$$

By observing that the latter distribution can actually be computed by a classical finite machine, we finish the proof.

5.5.2. Precise Definition of the Distributions $\mathcal{D}_{\mathcal{W}(\text{Pic}_K^0)}$, $\mathcal{D}_{\ddot{\mathcal{W}}(\text{Pic}_K^0)}$ and $\ddot{\mathcal{D}}_{\ddot{\mathcal{W}}(\text{Pic}_K^0)}$

Before defining the three relevant distributions, we first need to define the discretization of the Gaussian (in the random walk procedure) and of the circle (in the distribution procedure). The discretization of the continuous Gaussian happens by sampling a discrete Gaussian on a square grid of the log-unit hyperplane and the discretization of the hypercircle happens by taking equidistant points on this hypercircle.

Definition 5.12 (Orthogonal lattice in the log-unit hyperplane H). *By choosing an orthonormal basis $(\mathbf{b}_1, \dots, \mathbf{b}_r)$ of the r -dimensional vector space $H = \{(x_\sigma)_\sigma \in \log K_{\mathbb{R}} \mid \sum_\sigma x_\sigma = 0\}$, we define $\mathbb{Z}_H = \mathbf{b}_1 \mathbb{Z} + \dots + \mathbf{b}_r \mathbb{Z}$.*

The actual choice of the orthonormal basis doesn't matter in the proofs, so we will just work with the lattice $\mathbb{Z}_H \subseteq H$ and leaving the basis choice implicit. For $D \in \mathbb{N}_{>0}$, we denote $\frac{1}{D}\mathbb{Z}_H$ for the scaling of \mathbb{Z}_H by $\frac{1}{D}$, i.e., $\frac{1}{D}\mathbb{Z}_H = \frac{1}{D} \cdot (\mathbf{b}_1\mathbb{Z} + \dots + \mathbf{b}_r\mathbb{Z})$. To make the random walk procedure efficiently computable on a finite machine, we discretize the continuous Gaussian walk over H by sampling from a discrete Gaussian over $\frac{1}{D}\mathbb{Z}_H$.

Definition 5.13 (Sampling in the finite set $\ddot{\mathcal{C}}_M \subseteq \mathcal{C}_M \in K_{\mathbb{R}}$). *For a small discretization parameter $\varepsilon > 0$, we put $k = \sqrt{n} \cdot \lceil 1/\varepsilon \rceil$,*

$$\ddot{\mathcal{C}}_M^{(\varepsilon)} = \{(x_\sigma)_{\sigma} \in \mathcal{C}_M \mid x_\sigma = \pm M e^{2\pi i j/k} \text{ for some } j \in \mathbb{N}\}.$$

Recall that for real embeddings σ we have $x_\sigma = \pm M$, and for complex embeddings $x_{\bar{\sigma}} = \overline{x_\sigma}$, due to the fact that $\mathcal{C}_M \subseteq K_{\mathbb{R}}$. We often suppress the notation of ε in $\ddot{\mathcal{C}}_M$.

For most purposes, the precise definition of $\ddot{\mathcal{C}}_M^{(\varepsilon)}$ is not so important; what matters more is the fact that any point in \mathcal{C}_M is ε -close to $\ddot{\mathcal{C}}_M^{(\varepsilon)}$ (see Figure 5.3).

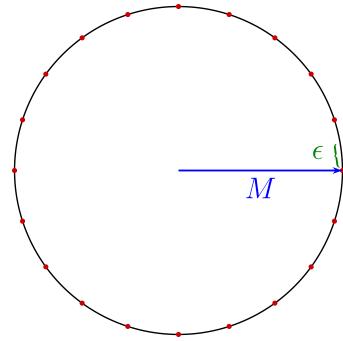


Figure 5.3.: Any point on the circle \mathcal{C}_M is ε -close to the red discretized circle $\ddot{\mathcal{C}}_M$.

Now we are ready to rigorously define the three distributions involved. We start with the distribution involving a continuous Gaussian and a continuous circle, Definition 5.14. The algorithm associated with this distribution is Algorithm 4, with Algorithm 3 as a subroutine.

Then we proceed by the definition of a intermediate distribution, which has a discrete Gaussian sampling in the random walk procedure, but still has a continuous sample from the circle in the distribution procedure, see Definition 5.15. The difference between these two distributions is marked with the color blue. The algorithm associated with this distribution is Algorithm 6, still with Algorithm 3 as a subroutine. Also in this algorithm description, the differences are marked in the color blue.

The last distribution is the one that can be run on a finite computer and has both a discretized Gaussian and a discretized circle, see Definition 5.16. The difference between this distribution and the intermediate distribution is marked with the color green. The algorithm associated with this distribution is Algorithm 6, with the discrete sampling on the circle Algorithm 5 as a subroutine, where the differences with the original (Algorithm 3) is marked with the color green as well.

Definition 5.14 (Continuous Gaussian, continuous circle). *Denoting $P_B = \{\mathfrak{p} \text{ prime ideal of } \mathcal{O}_K \mid \mathcal{N}(\mathfrak{p}) \leq B\}$, the output distribution $\mathcal{D}_{\mathcal{W}(\text{Pic}_K^0)}$ of Algorithm 4 can be described by the following rule, for any integral ideal $\mathfrak{d} \in \mathcal{I}_K$.*

$$\mathcal{D}_{\mathcal{W}(\text{Pic}_K^0)}[\mathfrak{d}^{-1}] = \frac{1}{|P_B| \cdot \text{Vol}(\mathcal{C}_M)} \int_{c \in \mathcal{C}_M} \sum_{\mathfrak{p} \in P_B} \int_{y \in H} \sum_{\substack{v \in \mathfrak{p} \\ (v) = \mathfrak{p}\mathfrak{d}}} \frac{\rho_s(e^y v / p - c)}{\rho_s(e^y \mathfrak{p} / p - c)} s^{-\mathfrak{r}} \rho_s(y) dy,$$

where $p = \mathcal{N}(\mathfrak{p})^{-1/n}$.

Definition 5.15 (Discrete Gaussian, continuous circle). *Denoting $P_B = \{\mathfrak{p} \text{ prime ideal of } \mathcal{O}_K \mid \mathcal{N}(\mathfrak{p}) \leq B\}$, the output distribution $\mathcal{D}_{\tilde{\mathcal{W}}(\text{Pic}_K^0)}$, where the continuous Gaussian $\mathcal{G}_{H,s}$ in Algorithm 4 is replaced by a discrete Gaussian $\mathcal{G}_{\frac{1}{D}\mathbb{Z}_H,s}$, can be described by the following rule, for any integral ideal $\mathfrak{d} \in \mathcal{I}_K$.*

$$\begin{aligned} \mathcal{D}_{\tilde{\mathcal{W}}(\text{Pic}_K^0)}[\mathfrak{d}^{-1}] &= \\ \frac{1}{|P_B| \cdot \text{Vol}(\mathcal{C}_M)} \int_{c \in \mathcal{C}_M} \sum_{\mathfrak{p} \in P_B} \sum_{\tilde{y} \in \frac{1}{D}\mathbb{Z}_H} \sum_{\substack{v \in \mathfrak{p} \\ (v) = \mathfrak{p}\mathfrak{d}}} & \frac{\rho_s(\lceil e^{\tilde{y}} \rceil v / p - c)}{\rho_s(\lceil e^{\tilde{y}} \rceil \mathfrak{p} / p - c)} \frac{\rho_s(\tilde{y})}{\rho_s(\frac{1}{D}\mathbb{Z}_H)}, \end{aligned}$$

where $\lceil e^{\tilde{y}} \rceil$ means that $e^{\tilde{y}}$ is computed with $\lceil \log_2 D \rceil$ bits of precision in all coordinates, and where $p = \mathcal{N}(\mathfrak{p})^{-1/n}$.

Definition 5.16 (Discrete Gaussian, discrete circle). Denoting $P_B = \{\mathfrak{p} \text{ prime ideal of } \mathcal{O}_K \mid \mathcal{N}(\mathfrak{p}) \leq B\}$, the output distribution $\mathcal{D}_{\tilde{\mathcal{W}}(\text{Pic}_K^0)}$, where the continuous Gaussian $\mathcal{G}_{H,s}$ in Algorithm 4 is replaced by a discrete Gaussian $\mathcal{G}_{\frac{1}{D}\mathbb{Z}_H, s}$, and the continuous uniform distribution on \mathcal{C}_M in $\mathcal{D} : \text{Pic}_K^0 \rightarrow L_1(\mathcal{I}_K)$ is replaced by a uniform distribution over the finite set $\tilde{\mathcal{C}}_M$ can be described by the following rule, for any integral ideal $\mathfrak{d} \in \mathcal{I}_K$.

$$\tilde{\mathcal{D}}_{\tilde{\mathcal{W}}(\text{Pic}_K^0)}[\mathfrak{d}^{-1}] = \frac{1}{|P_B| \cdot |\tilde{\mathcal{C}}_M|} \sum_{\tilde{c} \in \tilde{\mathcal{C}}_M} \sum_{\mathfrak{p} \in P_B} \sum_{\tilde{y} \in \frac{1}{D}\mathbb{Z}_H} \sum_{\substack{v \in \mathfrak{p} \\ (v) = \mathfrak{p}\mathfrak{d}}} \frac{\rho_s(\lceil e^{\tilde{y}} \rceil v/p - \tilde{c})}{\rho_s(\lceil e^{\tilde{y}} \rceil \mathfrak{p}/p - \tilde{c})} \frac{\rho_s(\tilde{y})}{\rho_s(\frac{1}{D}\mathbb{Z}_H)},$$

where $\lceil e^{\tilde{y}} \rceil$ means that $e^{\tilde{y}}$ is computed with $\lceil \log_2 D \rceil$ bits of precision in all coordinates, and where $p = \mathcal{N}(\mathfrak{p})^{-1/n}$.

5.5.3. Discretized Algorithm Analogues

In the following text we treat the discrete analogues of Algorithm 4 and Algorithm 3. We show that these discretized algorithms (Algorithm 6 and Algorithm 5) run in polynomial time with respect to the input size and that their output distribution does not differ significantly from their continuous counterparts.

We start with defining the algorithms and showing that they run in polynomial time. The remainder of this chapter, Section 5.5.4, is devoted to showing that the discretized and non-discretized algorithms indeed yield almost the same distribution.

Lemma 5.17. Algorithm 5 is correct and runs within time polynomial in $\log |\Delta_K|$, $\text{size}(M_{\mathfrak{a}})$, $\log(1/\varepsilon)$ and $\text{size}(x)$.

Proof. The input of this algorithm is given by the vector $x \in K_{\mathbb{R}}$ (given in a finite precision representation) and a basis matrix $B_{\mathfrak{a}}$ of the ideal \mathfrak{a} .

Algorithm 5: Sampling efficiently from a distribution very close to $\mathcal{D}_{x\mathfrak{a}}$, discretized

Require: A pair $(x, \mathfrak{a}) \in K_{\mathbb{R}} \times \mathcal{I}_K$ such that $\mathcal{N}(\mathfrak{a}) \prod_{\sigma} x_{\sigma} = 1$.

Ensure: An sample from a distribution (14ϵ) -close to the distribution $\mathcal{D}_{x\mathfrak{a}}$ in the total variation distance.

- 1: Put $\varsigma = 2^{n+1} \cdot n \cdot |\Delta_K|$ and $M = 2\sqrt{n} \cdot \varsigma$.
- 2: Sample a center $\ddot{c} = (\ddot{c}_{\sigma})_{\sigma}$ uniformly in the finite subset $\ddot{\mathcal{C}}_M := \ddot{\mathcal{C}}_M^{(\epsilon/n)} \subseteq \mathcal{C}_M = \{(y_{\sigma})_{\sigma} \mid |y_{\sigma}| = M \text{ for all embeddings } \sigma\}$. Where $\ddot{\mathcal{C}}_M$ is such that any point in \mathcal{C}_M is ϵ/n -close to $\ddot{\mathcal{C}}_M$ (see Definition 5.13)
- 3: Sample from the discrete Gaussian $\mathcal{G}_{x\mathfrak{a}, \varsigma, \ddot{c}}$ with respect to the ideal lattice $x\mathfrak{a}$ with center $\ddot{c} = (\ddot{c}_{\sigma})_{\sigma}$ and standard deviation ς , leading to some $v \in x\mathfrak{a}$.
- 4: **return** the inverse integral ideal $\mathfrak{d}^{-1} = v^{-1}x\mathfrak{a} \in \mathcal{I}_K$

We denote with $\text{size}(x)$ the number of bits needed to represent $x \in K_{\mathbb{R}}$ and with $\text{size}(B_{\mathfrak{a}})$ the number of bits needed to represent the basis $B_{\mathfrak{a}}$ (see Section 2.1).

We go through the lines of Algorithm 5 to examine the running time. Line 1 can clearly be done in linear time in $\log(|\Delta_K|)$ and n . Line 2 samples from in set $\ddot{\mathcal{C}}_M^{(\epsilon)}$, which are essentially at most⁵ $n/2$ independent samples of the discretized circle $\{Me^{2\pi ij/D} \mid j \in \mathbb{N}\}$, with $D = \sqrt{n}M\lceil 1/\epsilon \rceil$. One such sample takes time linear in $\log M = O(\log |\Delta_K|)$ and $\log(1/\epsilon)$, so a sample from $\ddot{\mathcal{C}}_M^{(\epsilon)}$ costs $O(n(\log |\Delta_K| + \log(1/\epsilon)))$. Line 3 uses Klein's algorithm [GPV08; Kle00] to sample from the discrete Gaussian $\mathcal{G}_{x\mathfrak{a}, \varsigma, \ddot{c}}$, which runs in time polynomial in $\text{size}(B_{\mathfrak{a}})$ and $\text{size}(x)$, by an adaptation of [GPV08, Thm. 4.1] for an exponentially small statistical distance. An additional property of Klein's algorithm is that the output $v \leftarrow \mathcal{G}_{x\mathfrak{a}, \varsigma, \ddot{c}}$ is actually polynomially bounded by $\text{size}(x)$ and $\text{size}(B_{\mathfrak{a}})$. The last line, line 4, uses ideal division and multiplication, which (naively) takes the time to solve a system of equations involving a $n^2 \times n^2$ matrix (see [Coh99, §4.8.4]) having

⁵At most half of n , because of the complex conjugate embeddings

entry sizes comparable to that of $\text{size}(B_{\mathfrak{a}})$ and $\text{size}(x)$; therefore this can be done within polynomial time in $\log |\Delta_K|$, $\text{size}(B_{\mathfrak{a}})$ and $\text{size}(x)$. As all lines can be computed in polynomial time of $\text{size}(B_{\mathfrak{a}})$, $\text{size}(x)$ and $\log |\Delta_K|$, the result follows.

The correctness is proven later, in Lemma 5.20. □

Lemma 5.18. *Algorithm 6 is correct and runs within time polynomial in $\log |\Delta_K|$, $\text{size}(M_{\mathfrak{a}})$, $\log[\Lambda_K : C]$ and $\text{size}(x)$, and uses one call to a γ -Hermite SVP oracle.*

Proof. The input of this algorithm is given by the vector $x \in K_{\mathbb{R}}$ (given in a finite precision representation) and a basis matrix $B_{\mathfrak{a}}$ of the ideal \mathfrak{a} . We denote with $\text{size}(x)$ the number of bits needed to represent $x \in K_{\mathbb{R}}$ and with $\text{size}(B_{\mathfrak{a}})$ the number of bits needed to represent the basis $B_{\mathfrak{a}}$.

Since $\log |\text{Pic}_K^0| = O(\log |\Delta_K|)$ (see Lemma 2.17), $n = O(\log |\Delta_K|)$ and $\eta_1(C^*) \leq \eta_1(\Lambda_K^*) \leq \tilde{O}(n)$ (see the proof of Proposition 5.10) the quantity $\log B$ is polynomially bounded in $\log |\Delta_K|$ and $\log[\Lambda_K : C]$. Similarly, $\log D$, the logarithm of the discretization parameter of the Gaussian, is polynomially bounded by $\log |\Delta_K|$ and $\log(\varepsilon^{-1})$.

We go through all steps of Algorithm 6 to estimate the running time. Step 1 of Algorithm 6 runs within time quasi-linear in $\log B$. Step 2 involves the sampling a random prime ideal \mathfrak{p} and the multiplication of ideals \mathfrak{a} and \mathfrak{p} . The random sampling can be done within polynomial time (see Lemma 2.14). The product $\mathfrak{p}\mathfrak{a}$ can be computed by reducing the $n^2 \times n$ matrix consisting of the products of the respective \mathbb{Z} -generators of \mathfrak{a} and \mathfrak{p} which runs in time polynomial in n , $\text{size}(\mathfrak{a})$ and $\log B$ (where B is the maximum norm of \mathfrak{p}). Step 3 consists of discrete Gaussian sampling in the lattice $\frac{1}{D}\mathbb{Z}_H$ with standard deviation s satisfying $\tilde{O}(n) \leq 1/s \leq \log(n)^2$. An adaptation of [GPV08, Thm. 4.1] shows that this can be done in time polynomially bounded by $\log D$ and n , i.e. bounded by $\log |\Delta_K|$ and $\log(\varepsilon^{-1})$. An additional property of this sampling is that the output is polynomially bounded as well. Step 4 is just rescaling, which has no serious impact on

Algorithm 6: The worst-case to average-case reduction algorithm, discretized

Require:

- A pair $(x, \mathfrak{a}) \in K_{\mathbb{R}} \times \mathcal{I}_K$ satisfying $\mathcal{N}(\mathfrak{a}) \prod_{\sigma} x_{\sigma} = 1$.
- The values $[\Lambda_K : C]$ and $\eta_1(C^*)$ of a suitable sublattice $C \subseteq \Lambda_K$ of the logarithmic unit lattice,
- An oracle \mathcal{A} that solves γ -Hermite SVP in \mathfrak{d}^{-1} whenever $\mathfrak{d}^{-1} \leftarrow \mathcal{D}_{\mathcal{U}(\text{Pic}_K^0)}$.
- An error parameter $\varepsilon > 0$

Ensure: A vector $\alpha \in x\mathfrak{a}$ that is a solution to $B^{1/n} \cdot \gamma$ -Hermite SVP in the ideal lattice $x\mathfrak{a}$, i.e.,

$$\|\alpha\| \leq \gamma \cdot B^{1/n} \cdot \det(x\mathfrak{a})^{1/n},$$

where $B = \tilde{O}(4^n \cdot \eta_1(C^*)^n \cdot [\Lambda_K : C] \cdot |\text{Pic}_K^0| \cdot (\log |\Delta_K|)^2)$, or, *failure*.

- 1: Put $s = \max(\sqrt{2} \cdot \eta_1(C^*), (\log n)^2)$ and $B = \tilde{O}(4^n \cdot \eta_1(C^*)^n \cdot [\Lambda_K : C] \cdot |\text{Pic}_K^0| \cdot (\log |\Delta_K|)^2)$ as in Corollary 5.8.
- 2: Multiply the ideal \mathfrak{a} by a prime ideal \mathfrak{p} uniformly sampled from the set $\{\mathfrak{p} \text{ prime ideal of } \mathcal{O}_K \mid \mathcal{N}(\mathfrak{p}) \leq B\}$, yielding $\mathfrak{a}\mathfrak{p}$.
- 3: Sample $\ddot{y} \leftarrow \mathcal{G}_{s, \frac{1}{D}\mathbb{Z}_H}$, where $D = 2^{n+2} \cdot n^4 \cdot \lceil |\Delta_K| \cdot \varepsilon^{-1} \rceil$ and \mathbb{Z}_H is an orthonormal basis of the hyperplane H where the logarithmic unit lattice lives in (see Definition 5.12).
- 4: Put $p = \mathcal{N}(\mathfrak{p})^{1/n}$, so that $e^y x \mathfrak{a} \mathfrak{p} / p$ has norm 1, where $e^y \in K_{\mathbb{R}}$ is the component-wise exponentiation of $y \in H$.
- 5: Sample $\mathfrak{d}^{-1} \leftarrow \mathcal{D}_{(\lfloor e^{\ddot{y}} \rfloor \cdot x / p, \mathfrak{a}\mathfrak{p})}$ using Algorithm 5, where $\lfloor e^{\ddot{y}} \rfloor \in K_{\mathbb{R}}$ is the component-wise exponentiation of $\ddot{y} \in H$, computed with $\lceil \log_2 D \rceil$ bits of precision in all coordinates. Furthermore, let $v \in e^y x \mathfrak{a} \mathfrak{p} / p$ be the additional output of Algorithm 5 that satisfies $\mathfrak{d}^{-1} = v^{-1} \lfloor e^{\ddot{y}} \rfloor x \mathfrak{a} \mathfrak{p} / p$.
- 6: Invoke the γ -Hermite SVP oracle \mathcal{A} on $\mathcal{D}_{\mathcal{U}(\text{Pic}_K^0)}$ to find a $\kappa \in \mathfrak{d}^{-1}$ for which holds $\|\kappa\| \leq \gamma \cdot \det(\mathfrak{d}^{-1})^{1/n}$
- 7: **return** $p \cdot (\lfloor e^{\ddot{y}} \rfloor)^{-1} \cdot v \cdot \kappa \in x\mathfrak{a}$.

the running time. Step 5 uses Algorithm 5, which runs in time polynomially bounded by $\text{size}(\lfloor e^{\tilde{y}} \rfloor \cdot x/p)$, $\text{size}(M_{\mathfrak{a}})$ and $\log B$. As $\text{size}(\lfloor e^{\tilde{y}} \rfloor \cdot x/p)$ can be linearly bounded by $\log B$, $\text{size}(x)$ and $\log D$ (because $\lfloor e^{\tilde{y}} \rfloor$ is computed with relative bit precision $\log_2 D$), this step is polynomially bounded as well in $\log |\Delta_K|$, $\log \varepsilon^{-1}$ and $\text{size}(x)$. Step 6 invokes the γ -Hermite SVP oracle once. Step 7 just rescales the element $\kappa \in \mathfrak{d}^{-1}$ without a serious impact on the running time.

Later in this section we prove two closeness lemmas, namely Lemma 5.19 and Lemma 5.20. From those two lemmas, one obtains the desired closeness of distributions of the sampling mechanism of \mathfrak{d}^{-1} ; this proves the correctness.

□

5.5.4. Closeness Proofs

Sampling the Gaussian walk in a discrete manner doesn't spoil the resulting distribution

Lemma 5.19. *Let K be a number field and let $1 > s > 0$ be a given Gaussian spread parameter for the continuous part of the random walk, let $\varepsilon > 0$ be a given error parameter and let $M = 2 \cdot n^{3/2} \cdot 2^{n+1} \cdot |\Delta_K|$ as in Algorithm 5. Let $\frac{1}{D} \mathbb{Z}_H \subseteq H$ be the discretization of the Log-unit space to compute the discrete Gaussian analogue of the continuous part of the random walk, with $D \in \mathbb{N}$ such that $D \geq \lceil (4 \cdot s^{-2} \sqrt{n} + 100 \cdot n^2 M) \cdot 1/\varepsilon \rceil$.*

Then

$$\|\mathcal{D}_{\tilde{\mathcal{W}}(\text{Pic}_K^0)} - \mathcal{D}_{\mathcal{W}(\text{Pic}_K^0)}\|_1 \leq 18 \cdot \varepsilon$$

Proof. Examining the definitions of the distributions $\mathcal{D}_{\mathcal{W}(\text{Pic}_K^0)}$ and $\mathcal{D}_{\tilde{\mathcal{W}}(\text{Pic}_K^0)}$ (see Definitions 5.14 and 5.15), we can apply the triangle inequality and a

norm inequality, to directly deduce

$$\begin{aligned} & \| \mathcal{D}_{\mathcal{W}(\text{Pic}_K^0)} - \mathcal{D}_{\tilde{\mathcal{W}}(\text{Pic}_K^0)} \|_1 \\ & \leq \max_{c \in \mathcal{C}_M} \sum_{\mathfrak{p} \in P_B} \sum_{\mathfrak{d}} \underbrace{\sum_{\substack{v \in \mathfrak{p} \\ (v) = \mathfrak{p}\mathfrak{d}}} \left| \int_{y \in H} \frac{\rho_s(e^y v/p - c)}{\rho_s(e^y \mathfrak{p}/p - c)} s^{-\mathfrak{r}} \rho_s(y) dy - \sum_{\tilde{y} \in \frac{1}{D} \mathbb{Z}_H} \frac{\rho_s(\lfloor e^{\tilde{y}} \rfloor v/p - c)}{\rho_s(\lfloor e^{\tilde{y}} \rfloor \mathfrak{p}/p - c)} \frac{\rho_s(\tilde{y})}{\rho_s(\frac{1}{D} \mathbb{Z}_H)} \right| }_{\text{bracket}}. \end{aligned} \quad (5.70)$$

Therefore, we can focus on the quantity in the bracket of Equation (5.70) for a fixed prime ideal $\mathfrak{p} \in P_B$ and a fixed center $c \in \mathcal{C}_M$ from the M -circle. We rewrite the term within the absolute value signs by using a block tiling of the orthonormal lattice $\frac{1}{D} \mathbb{Z}_H \subseteq H$ (see Definition 5.12) with fundamental domain F_H satisfying $\text{Vol}(F_H) = D^{-\mathfrak{r}}$. Observing that we can collapse the summation $\sum_{\mathfrak{d}} \sum_{\substack{v \in \mathfrak{p} \\ (v) = \mathfrak{p}\mathfrak{d}}}$ to $\sum_{v \in \mathfrak{p}}$ (as the sum with \mathfrak{d} is over integral ideals), we obtain that the quantity in the bracket of Equation (5.70) is at most

$$\sum_{v \in \mathfrak{p}} \left| \sum_{\tilde{y} \in \frac{1}{D} \mathbb{Z}_H} \int_{y \in \tilde{y} + F_H} \underbrace{\frac{\rho_s(e^y v/p - c)}{\rho_s(e^y \mathfrak{p}/p - c)} s^{-\mathfrak{r}} \rho_s(y)}_A - \underbrace{\frac{\rho_s(\lfloor e^{\tilde{y}} \rfloor v/p - c)}{\rho_s(\lfloor e^{\tilde{y}} \rfloor \mathfrak{p}/p - c)} \frac{D^{\mathfrak{r}} \rho_s(\tilde{y})}{\rho_s(\frac{1}{D} \mathbb{Z}_H)}}_B dy \right|$$

Applying the triangle inequality, switching integrals and sums, and using the identity $AB - A'B' = B(A - A') + (B - B')A'$, above equation is at most

$$\sum_{\tilde{y} \in \frac{1}{D} \mathbb{Z}_H} \int_{y \in \tilde{y} + F_H} s^{-\mathfrak{r}} \rho_s(y) \underbrace{\sum_{v \in \mathfrak{p}} \left| \frac{\rho_s(e^y v/p - c)}{\rho_s(e^y \mathfrak{p}/p - c)} - \frac{\rho_s(\lfloor e^{\tilde{y}} \rfloor v/p - c)}{\rho_s(\lfloor e^{\tilde{y}} \rfloor \mathfrak{p}/p - c)} \right|}_{\| \mathcal{G}_{\mathfrak{p}/p, \varsigma/e^y, c} - \mathcal{G}_{\mathfrak{p}/p, \varsigma/\lfloor e^{\tilde{y}} \rfloor, c} \|} dy \quad (5.71)$$

$$+ \sum_{\tilde{y} \in \frac{1}{D} \mathbb{Z}_H} \int_{y \in \tilde{y} + F_H} \left| s^{-\mathfrak{r}} \rho_s(y) - \frac{D^{\mathfrak{r}} \rho_s(\tilde{y})}{\rho_s(\frac{1}{D} \mathbb{Z}_H)} \underbrace{\sum_{v \in \mathfrak{p}} \frac{\rho_s(e^{\tilde{y}} v/p - c)}{\rho_s(e^{\tilde{y}} \mathfrak{p}/p - c)}}_{=1} dy \right| \quad (5.72)$$

First part of the sum, Equation (5.71). Apply Lemma A.39 to show that the two Gaussians are reasonably close to each other. Writing $e^{\tilde{y}} = \lfloor e^{\tilde{y}} \rfloor$, we have $\|\tilde{y} - \tilde{y}\| \leq \|e^{\tilde{y}} - \lfloor e^{\tilde{y}} \rfloor - 1\| \leq \frac{\sqrt{n}}{D}$ because $e^{\tilde{y}}$ is the $\log_2(D)$ -bit precision

relative approximation of $e^{\ddot{y}}$. By construction, we have $\|y - \ddot{y}\| \leq \frac{\sqrt{n}}{D}$ as well, because $y \in \ddot{y} + F_H$, therefore,

$$\|y - \ddot{y}\| \leq \frac{2\sqrt{n}}{D}.$$

Because $\varsigma > \eta_\varepsilon(x\mathfrak{a})$ for all ideal lattices $x\mathfrak{a} \in \text{IdLat}_K$, we can apply Lemma A.39 with $\delta = \frac{2\sqrt{n}}{D}$. Since $M > \varsigma > 1$, $\|c\| = \sqrt{n} \cdot M$ (because $c \in \mathcal{C}_M$) and $D \geq 100 \cdot M \cdot n^2/\varepsilon$, we obtain

$$\begin{aligned} \|\mathcal{G}_{\mathfrak{p}, \varsigma/e^y, c} - \mathcal{G}_{\mathfrak{p}, \varsigma/\lfloor e^{\ddot{y}} \rfloor, c}\| &\leq 8\varepsilon + 4\pi\left(\frac{1}{\varsigma^2} + n + 2n\|c\|\right) \cdot \|y - \ddot{y}\| \\ &\leq 8\varepsilon + \frac{100 \cdot M \cdot n^2}{D} = 9 \cdot \varepsilon. \end{aligned} \quad (5.73)$$

Since this bound is independent of $y \in H$ and $\ddot{y} \in \frac{1}{D}\mathbb{Z}_H$, and since

$$\sum_{\ddot{y} \in \frac{1}{D}\mathbb{Z}_H} \int_{y \in \ddot{y} + F_H} s^{-r} \rho_s(y) dy = 1,$$

we deduce that Equation (5.71) must also be bounded by $9 \cdot \varepsilon$.

Second part of the sum, Equation (5.72). One can apply smoothing arguments; since $s < 1$, we have $s \geq s^2 \geq \frac{\sqrt{n} \cdot \varepsilon^{-1}}{D} \geq \frac{\log(2n(1+\varepsilon^{-1}))}{D} \geq \log(2n(1+\varepsilon^{-1}))\lambda_r(\frac{1}{D}\mathbb{Z}_H) \geq \eta_\varepsilon(\frac{1}{D}\mathbb{Z}_H)$ (see [MR07, Lm. 3.3]). Therefore,

$$s^{-r} \in (1 - 2\varepsilon, 1 + 2\varepsilon) \cdot \frac{D^r}{\rho_s(\frac{1}{D}\mathbb{Z}_H + y)} \quad \text{for all } y \in H.$$

Putting this into Equation (5.72), we obtain, using Lemma A.37, using the lower bound on D and $\text{Vol}(F_H) = D^{-r}$,

$$\begin{aligned} &\sum_{\ddot{y} \in \frac{1}{D}\mathbb{Z}_H} \int_{y \in F_H} \left| s^{-r} \rho_s(\ddot{y} + y) - \frac{D^r \rho_s(\ddot{y})}{\rho_s(\frac{1}{D}\mathbb{Z}_H)} \right| dy \\ &\leq 4\varepsilon + \max_{y \in F_H} \sum_{\ddot{y} \in \frac{1}{D}\mathbb{Z}_H} \underbrace{\left| \frac{\rho_s(\ddot{y} + y)}{\rho_s(\frac{1}{D}\mathbb{Z}_H + y)} - \frac{\rho_s(\ddot{y})}{\rho_s(\frac{1}{D}\mathbb{Z}_H)} \right|}_{\|D_{\frac{1}{D}\mathbb{Z}_H, s, y} - D_{\frac{1}{D}\mathbb{Z}_H, s, 0}\|} \\ &\leq 8\varepsilon + \left(\frac{\pi}{s^2} + 2\pi n \right) \max_{y \in F_H} \|y\| \leq 8\varepsilon + \left(\frac{\pi}{s^2} + 2\pi n \right) \frac{\sqrt{n}}{D} \leq 9 \cdot \varepsilon \end{aligned} \quad (5.74)$$

Combining the upper bound on the first and the second part of the sum (see Equations (5.73) and (5.74)), we obtain the result. \square

The difference between $\mathcal{D}_{\mathcal{P}(\mathrm{Pic}_K^0)}$ and $\ddot{\mathcal{D}}_{\mathcal{P}(\mathrm{Pic}_K^0)}$, the one with a discretized circle, is negligible for all distributions \mathcal{P}

Lemma 5.20. Let K be a number field and let $1 > \varepsilon > 0$ be a given error parameter. Let $\ddot{\mathcal{C}}_M \subseteq \mathcal{C}_M$ be a discretization of \mathcal{C}_M as in Definition 5.13. Let furthermore $\mathcal{P} \in L_1(\mathrm{Pic}_K^0)$ be any distribution on Pic_K^0 (i.e., $\int_{\mathbf{a} \in \mathrm{Pic}_K^0} \mathcal{P}(\mathbf{a}) d\mathbf{a} = 1$). Then we have

$$\|\mathcal{D}_{\mathcal{P}(\mathrm{Pic}_K^0)} - \ddot{\mathcal{D}}_{\mathcal{P}(\mathrm{Pic}_K^0)}\| \leq 14\varepsilon$$

Proof. The definitions of the two distributions read as follows, for integral ideals $\mathfrak{d} \in \mathcal{I}_K$.

$$\mathcal{D}_{\mathcal{P}(\mathrm{Pic}_K^0)}[\mathfrak{d}^{-1}] = \int_{c \in \mathcal{C}_M} \int_{\mathbf{a} \in \mathrm{Pic}_K^0} \mathcal{P}(\mathbf{a}) \sum_{\substack{v \in \mathbf{a} \\ (v) = \mathbf{a}\mathfrak{d}}} \frac{\rho_\varsigma(v - c)}{\rho_\varsigma(\mathbf{a} - c)} d\mathbf{a} dc$$

$$\ddot{\mathcal{D}}_{\mathcal{P}(\mathrm{Pic}_K^0)}[\mathfrak{d}^{-1}] = \frac{1}{|\ddot{\mathcal{C}}_M|} \sum_{\ddot{c} \in \ddot{\mathcal{C}}_M} \int_{\mathbf{a} \in \mathrm{Pic}_K^0} \mathcal{P}(\mathbf{a}) \sum_{\substack{v \in \mathbf{a} \\ (v) = \mathbf{a}\mathfrak{d}}} \frac{\rho_\varsigma(v - \ddot{c})}{\rho_\varsigma(\mathbf{a} - \ddot{c})} d\mathbf{a}$$

By grouping integrals and summation signs, and splitting up the integral over \mathcal{C}_M over multiple ‘arcs’ $A_{\ddot{c}}$ for $\ddot{c} \in \ddot{\mathcal{C}}_M$ (that satisfy $\|c - \ddot{c}\| < \varepsilon/n$ for all $c \in A_{\ddot{c}}$), we obtain

$$\begin{aligned} & \|\mathcal{D}_{\mathcal{P}(\mathrm{Pic}_K^0)} - \ddot{\mathcal{D}}_{\mathcal{P}(\mathrm{Pic}_K^0)}\| \\ &= \sum_{\mathfrak{d} \in \mathcal{I}_K} \left| \int_{\mathbf{a} \in \mathrm{Pic}_K^0} \mathcal{P}(\mathbf{a}) \sum_{\substack{v \in \mathbf{a} \\ (v) = \mathbf{a}\mathfrak{d}}} \left(\int_{c \in \mathcal{C}_M} \frac{\rho_\varsigma(v - c)}{\rho_\varsigma(\mathbf{a} - c)} dc - \frac{1}{|\ddot{\mathcal{C}}_M|} \sum_{\ddot{c} \in \ddot{\mathcal{C}}_M} \frac{\rho_\varsigma(v - \ddot{c})}{\rho_\varsigma(\mathbf{a} - \ddot{c})} \right) d\mathbf{a} \right| \\ &= \sum_{\mathfrak{d} \in \mathcal{I}_K} \left| \int_{\mathbf{a} \in \mathrm{Pic}_K^0} \mathcal{P}(\mathbf{a}) \sum_{\substack{v \in \mathbf{a} \\ (v) = \mathbf{a}\mathfrak{d}}} \left(\sum_{\ddot{c} \in \ddot{\mathcal{C}}_M} \int_{c \in A_{\ddot{c}}} \left(\frac{\rho_\varsigma(v - c)}{\rho_\varsigma(\mathbf{a} - c)} - \frac{\rho_\varsigma(v - \ddot{c})}{\rho_\varsigma(\mathbf{a} - \ddot{c})} \right) dc \right) d\mathbf{a} \right|. \end{aligned} \tag{5.75}$$

Applying the triangle inequality, switching integral and summation signs appropriately, collapsing the summation $\sum_{\mathfrak{d} \in \mathcal{I}_K} \sum_{\substack{v \in \mathbf{a} \\ (v) = \mathbf{a}\mathfrak{d}}}$ to $\sum_{v \in \mathbf{a}}$ (as \mathfrak{d} ranges

over integral ideals) and replacing the integral over Pic_K^0 by the maximum, we obtain that Equation (5.75) must be bounded by

$$\begin{aligned} \max_{\mathbf{a} \in \text{Pic}_K^0} \left(\sum_{\ddot{c} \in \mathcal{C}_M} \int_{c \in A_{\ddot{c}}} \underbrace{\sum_{v \in \mathbf{a}} \left| \frac{\rho_{\varsigma}(v - c)}{\rho_{\varsigma}(\mathbf{a} - c)} - \frac{\rho_{\varsigma}(v - \ddot{c})}{\rho_{\varsigma}(\mathbf{a} - \ddot{c})} \right| dc}_{\|D_{\mathbf{a}, \varsigma, c} - D_{\mathbf{a}, \varsigma, \ddot{c}}\|} \right) &\leq (4 + \frac{\pi}{\varsigma^2} + 2\pi n)\varepsilon/n \\ &\leq 14\varepsilon. \end{aligned} \quad (5.76)$$

This holds by the fact that $\|c - \ddot{c}\| < \varepsilon/n$ and $\varsigma > 1$, together with Lemma A.37, which bounds the total variation distance between two discrete Gaussians with different centers. \square

Conclusion

Applying Lemma 5.19 and Lemma 5.20 with $\mathcal{P}(\text{Pic}_K^0) = \ddot{\mathcal{W}}(\text{Pic}_K^0)$, and using Algorithm 6 for the running time, we obtain Theorem 5.11.

6. Ideal sampling

6.1. Summary

Many algorithms in cryptography and algorithmic number theory rely on finding elements α in an ideal \mathfrak{a} such that their quotient $\alpha\mathfrak{a}^{-1}$ is easy to factor (e.g., prime, near-prime or B -smooth). Such algorithms are typically analyzed only heuristically, by treating $\alpha\mathfrak{a}^{-1}$ as a uniform ideal, and applying density results for the sets of prime ideals or smooth ideals. The result of this chapter allows to adjust this strategy and make the reasoning rigorous.

The beginning of this chapter is devoted to showing that, for an ideal \mathfrak{a} that is uniformly distributed in the Arakelov class group, one can rigorously analyze the probability of $\alpha\mathfrak{a}^{-1}$ being in a certain ideal set (e.g., the prime ideals or smooth ideals). This probability can be shown to be very much related to the *density* of the ideal set involved, a notion from analytic number theory.

In the later part of this chapter we invoke the random walk theorem from Chapter 4, which allows to randomize any fixed ideal \mathfrak{a} into a randomly distributed ideal $\tilde{\mathfrak{a}}$ in the Arakelov class group. This *randomized* ideal can then be used to sample an $\alpha \in \tilde{\mathfrak{a}}$ from, with a rigorous probability. Sampling α from $\tilde{\mathfrak{a}}$ instead of \mathfrak{a} does not affect the usefulness of α , since the randomization – apart from a small distortion – happens only by multiplying \mathfrak{a} with small prime ideals. I.e., the quotient $\alpha\mathfrak{a}^{-1}$ only differs from $\alpha\tilde{\mathfrak{a}}^{-1}$ by small prime factors, meaning that if the one is easy to factor, the other is as well.

6.2. Introduction

In this chapter, we apply the random walk theorem of Chapter 4 to tackle the following problem that arises in multiple number-theoretic contexts [BF14; BP17; Buc88]. Let K be a number field, of degree n and discriminant Δ_K . Given an ideal $\mathfrak{a} \subset K$, sample an element $\alpha \in \mathfrak{a}$ such that the ideal $\alpha\mathfrak{a}^{-1}$ is easy to factor. In some cases (e.g., [BF14; Buc88; LL+93]), the fraction $\alpha\mathfrak{a}^{-1}$ is required to only have small prime factors, whereas in other cases (e.g., [BP17]), the fraction $\alpha\mathfrak{a}^{-1}$ is required to be a near-prime (i.e., at most one of its prime factors is allowed to be large).

In the literature and computer algebra systems (e.g., [CS08, §6.5] [BCP97; PAR19]), this task is performed by computing a reasonably short basis of the ideal \mathfrak{a} (by means of LLL, for example) and repeatedly randomly sampling reasonably short elements $\alpha \in \mathfrak{a}$ using this basis, until $\alpha\mathfrak{a}^{-1}$ is of the desired form. Assuming heuristically that the ideals $\alpha\mathfrak{a}^{-1}$ are more or less randomly distributed among ideals of bounded norm, one can use specific density results for subsets of ideals to obtain a heuristic estimate for the success probability of this method.

Even though the above approach appears to work in many practical cases, it is generally hard to prove anything in the direction of a rigorous lower bound for the success probability. A first obstacle is that the ideal $\alpha\mathfrak{a}^{-1}$ is not ‘random enough’ as, for example, it always lies in the ideal class $[\mathfrak{a}]^{-1}$. Even for principal ideal domains, a second obstacle is that the number of generators of (α) may vary unpredictably among sub-ideals (α) of \mathfrak{a} , resulting in some sub-ideals of \mathfrak{a} to be sampled more often than others, making the distribution of $\alpha\mathfrak{a}^{-1}$ skewed.

6.2.1. Our Technique

To resolve these issues, we slightly modify both the ideal \mathfrak{a} and the way $\alpha \in \mathfrak{a}$ is sampled. More precisely, the ideal \mathfrak{a} is replaced by $\tilde{\mathfrak{a}} = \mathfrak{a} \cdot \prod_i \mathfrak{p}_i$, where each \mathfrak{p}_i is a small, random prime ideal. The element α is then sampled

uniformly in the ideal $\tilde{\mathfrak{a}}$ intersected with a ‘distorted box’ in the canonical embedding space $K_{\mathbb{R}}$. More specifically, in the case of a totally real number field, the box is chosen as $\mathcal{N}(\tilde{\mathfrak{a}})^{1/n} \cdot B_{r,x} = \mathcal{N}(\tilde{\mathfrak{a}})^{1/n} \cdot \prod_j [-re^{x_j}, re^{x_j}] \subseteq K_{\mathbb{R}}$ with large enough $r > 0$, where $x_j \in \mathbb{R}$ satisfy $\sum_j x_j = 0$ and are distributed according to a Gaussian distribution.

This procedure mimics a random walk in the Arakelov class group, where multiplying by small primes accounts for the randomization at the finite places, whereas the distortion of the sampling boxes accounts for the randomization at the infinite places.

The idea behind this procedure is that, while it is hard to predict exactly how many generators of the ideal (α) are in $\mathfrak{a} \cap \mathcal{N}(\mathfrak{a})^{1/n} \cdot B_{r,0}$, the average number of such generators in $\tilde{\mathfrak{a}} \cap \mathcal{N}(\tilde{\mathfrak{a}})^{1/n} \cdot B_{r,x}$ is accurately predictable whenever $\tilde{\mathfrak{a}}$ and x are adequately randomized. Indeed, this quantity is given by the number of points of a shifted Log-unit lattice, intersected with a simplex; this number of points is hard to estimate for a given shift of the Log-unit lattice, but it is predictable on average, according to the following fact.

Lemma 6.1. *Let $S \subset \mathbb{R}^n$ be a measurable set and $\Lambda \subset \mathbb{R}^n$ a full rank lattice. Then, for a uniformly chosen $c \in \mathbb{R}^n / \Lambda$ it holds that $\mathbb{E}[|(\Lambda + c) \cap S|] = \text{Vol}(S) / \text{Vol}(\Lambda)$.*

Algorithmically, sampling uniformly in a box $\mathcal{N}(\mathfrak{a})^{1/n} \cdot B_{r,x}$ and element of an ideal \mathfrak{a} can be done in polynomial time with an LLL reduced basis, whenever $\log r = \text{poly}(n, \log |\Delta_K|)$. One can also strengthen this reduction as in [BF14; Buc88] for other time-quality trade-offs. Denoting \mathcal{S}_B the set of B -smooth ideals, and $\delta_{\mathcal{S}}[t]$ the density of ideals of norm $\leq t$ which belong to a given set of ideals \mathcal{S} , our (slightly simplified) main result is the following.

Main theorem. *Let \mathcal{S} be any set of integral ideals. Assuming the Riemann hypothesis for Hecke L -functions, there exists some $B = \text{poly}(\log |\Delta_K|)$, such that Algorithm 7 outputs in time $\text{poly}(\log |\Delta_K|, \text{size}(\mathcal{N}(\mathfrak{a})))$ an element $\alpha \in \mathfrak{a}$ such that $(\alpha)/\mathfrak{a} \in \mathcal{S} \cdot \mathcal{S}_B$ with probability at least $\frac{1}{3}\delta_{\mathcal{S}}[r^n] - 2^{-n}$.*

Since $\mathcal{N}((\alpha)/\mathfrak{a}) \approx r^n$, we have therefore formalized the heuristic that element sampling probability matches ideal density, up to a loss of $1/3$ on the probability, and up to an extra smooth ideal in \mathcal{S}_B . Moreover, the original purpose, namely finding a $\alpha \in \mathfrak{a}$ such that $\alpha\mathfrak{a}^{-1}$ can be easily factored, is not spoiled.

Including the modulus \mathfrak{m}

The non-simplified main result of this chapter (see Theorem 6.9) involves a *modulus* $\mathfrak{m} \subseteq \mathcal{O}_K$, an ideal that is to be ‘avoided’ in the computations.

Specifically, the main theorem states the probability that $(\alpha)/\mathfrak{a}$ is in a given ideal set, *given* the fact that $\alpha \equiv \tau$ modulo \mathfrak{m} for some fixed given $\tau \in K^{\mathfrak{m}}$. This particular generalization is included because of its usefulness for the computation of the power residue symbol (see Chapter 7). One recovers the main theorem, as described in this introduction, by putting $\mathfrak{m} = \mathcal{O}_K$.

6.2.2. Applications

As a direct application, one can prove that sampling $\alpha \in \mathfrak{a}$ such that $\alpha\mathfrak{a}^{-1}$ is a near-prime can be done efficiently in cyclotomic fields. This proves that the ‘principalization step’ in the power residue symbol algorithm of the author of this PhD thesis [BP17, §5.2] runs in polynomial time in the special case of cyclotomic number fields, and more generally in any family of number fields with small Dedekind zeta residue ρ_K .

The most general version of the result of this chapter (Theorem 6.9), involving a modulus \mathfrak{m} has even farther consequences. It does not only allow to remove one specific heuristic ([BP17, §5.2]), but can actually be applied in order to prove that the *entire* (slightly modified) algorithm for the power residue symbol is efficient (see Chapter 7).

We also hope that our technique could be helpful in proving other heuristic algorithms such as the index calculus algorithms of [Buc88; BF14] (computing class groups and unit groups), though other obstacles are expected. Not

only does it require universal bounds on the density of B -smooth ideal in large-degree number field, one also needs to ensure sufficient independence of the obtained multiplicative relations. For the second obstacle, further randomization techniques could turn out useful.

6.2.3. Related Works

We note that the issues we mention above have been circumvented in special cases. Building on a result of Seysen [Sey87], Hafner and McCurley [HM89] gave a provable algorithm for computing class-group and unit group of imaginary quadratic fields. This algorithm involves a random walk in the class group, which is used to prove that one can find a B -smooth *principal* ideal relatively efficiently. The idea of performing a random walk in the class group was reused in the algorithms of Buchmann [Buc88] and Biasse and Fieker [BF14], in a heuristic way. Finally, we note also that Schoof [Sch08] rephrased Buchmann's algorithm in the terms of Arakelov theory, and we heavily borrowed from his formalism.

6.3. Preliminaries

An important concept that plays a large role throughout the entire proof of the main theorem is that of a *generator of an Arakelov ray divisor*. This can be seen as a generalization of a generator of a principal ideal $\mathfrak{a} \subseteq \mathcal{O}_K$, which is an $\alpha \in \mathcal{O}_K$ satisfying $(\alpha) = \mathfrak{a}$. Such a generator α of the ideal \mathfrak{a} is called τ -equivalent (with respect to a modulus \mathfrak{m}) if $\alpha \equiv \tau \pmod{\mathfrak{m}}$ (note that this definition only makes sense if \mathfrak{m} and \mathfrak{a} are coprime).

The generalization to Arakelov ray divisors is very similar. As we can see Arakelov ray divisors as ideal lattices $x\mathfrak{a}$, a generator of such divisor is just an element in $K_{\mathbb{R}}$ of the shape $x\alpha$, where α is a generator of \mathfrak{a} . Of course, if \mathfrak{a} is not a principal ideal, there are no such generators. The τ -equivalent generators are just those $x\alpha \in K_{\mathbb{R}}$ for which $\alpha \equiv \tau \pmod{\mathfrak{m}}$. The precise definition is as follows.

6. Ideal sampling

Definition 6.2 (Generators of an Arakelov ray divisor). *Let $\tau \in K^{\mathfrak{m}}$ and let $\mathbf{a} \in \text{Div}_{K^{\mathfrak{m}}}$ be an Arakelov ray divisor with an infinite part \mathbf{a}_∞ and a finite part \mathbf{a}_f (see Equation (2.13)). We define the set of τ -equivalent generators $\text{Exp}(\mathbf{a})_\tau^\times \subseteq K_{\mathbb{R}}$ of \mathbf{a} by the following rule*

$$\text{Exp}(\mathbf{a})_\tau^\times := \begin{cases} \text{Exp}(\mathbf{a}_\infty) \cdot (\kappa \cdot \mathcal{O}_K^\times \cap \tau K^{\mathfrak{m},1}) \subseteq \text{Exp}(\mathbf{a}) & \text{if } \text{Exp}(\mathbf{a}_f) = (\kappa) \\ & \text{for some } \kappa \in K^{\mathfrak{m}} \\ \emptyset & \text{otherwise} \end{cases}$$

Equivalently, we can write

$$\text{Exp}(\mathbf{a})_\tau^\times = \{\alpha \in \text{Exp}(\mathbf{a}) \mid (\text{Exp}(-\mathbf{a}_\infty) \cdot \alpha) = \text{Exp}(\mathbf{a}_f) \text{ and } \text{Exp}(-\mathbf{a}_\infty) \cdot \alpha \in \tau K^{\mathfrak{m},1}\}.$$

The following specialization of the random walk theorem of Chapter 4 is tailored to the purposes of this chapter. These purposes require both N and B to be polynomially small, and s to be rather small as well, to ease sampling in the log-normally distorted box. There is no specific reason we chose *this* particular instantiation below, except for concreteness.

Theorem 6.3 (Random walks on the Arakelov ray class group, ERH). *Let $n = [K : \mathbb{Q}] \geq 2$, $s = 1/(1000 \cdot n^2)$ and let $\varepsilon > 0$ be an error parameter. There exists a bound $B = \tilde{O}(n^4[\log \log(1/\varepsilon)]^2 + n^2[\log(|\Delta_K| \mathcal{N}(\mathfrak{m}))]^2)$ such that for $N = \lfloor 8n + \log \mathcal{N}(\mathfrak{m}) + \log |\Delta_K| + 2 \log(1/\varepsilon) \rfloor$ the random walk distribution $[\mathcal{W}(B, N, s)]$ on $\text{Pic}_{K^{\mathfrak{m}}}^0$ is ε -close to uniform in $L_1(\text{Pic}_{K^{\mathfrak{m}}}^0)$, i.e.,*

$$\|[\mathcal{W}(B, N, s)] - \mathcal{U}(\text{Pic}_{K^{\mathfrak{m}}}^0)\|_1 \leq \varepsilon.$$

Proof. This formulation of the random walk theorem is obtained by instantiating Theorem 4.3 of Chapter 4 with $C = \Lambda_{K^{\mathfrak{m},1}}$, $s = 1/n^2$ and $k = 1$.

In that case, $B = \tilde{O}(n^4[\log \log(1/\varepsilon)]^2 + n^2[\log(|\Delta_K| \mathcal{N}(\mathfrak{m}))]^2)$, by simply suppressing polylogarithmic factors. By using the bounds $\eta_1(\Lambda_{K^{\mathfrak{m},1}}^*) \leq \eta_1(\Lambda_K^*) \leq 2000 \cdot (\mathfrak{r} + 1) \cdot \log(\mathfrak{r})^3 \leq 1000 \cdot n^2 = s^{-1}$ (see the proof of Proposition 5.10),

$\log \text{Vol}(\text{Pic}_{K^m}^0) \leq \log(\mathcal{N}(m)) + \log \text{Vol}(\text{Pic}_{K^m}^0) \leq \log(\mathcal{N}(m)) + \log |\Delta_K|$ (see Lemma 2.17) and $r \leq n$ one obtains that Theorem 4.3 applies, with

$$\begin{aligned} N &= \lfloor 8n + \log \mathcal{N}(m) + \log |\Delta_K| + 2 \log(1/\varepsilon) \rfloor \\ &\geq \frac{1}{2 \log n} \cdot (n \log(1000n^2) + \log \mathcal{N}(m) + \log |\Delta_K| + 2 \log(1/\varepsilon) + 2) \\ &\geq \frac{1}{2k \log n} \cdot (r \cdot \log(1/\tilde{s}) + \log |\text{Pic}_{K^m}^0| + 2 \log(1/\varepsilon) + 2). \end{aligned}$$

□

Remark 6.4. One can simplify the bounds on B and N in the theorem above by putting $\varepsilon = 2^{-n}$. In that case $B = \tilde{O}(n^2 \cdot (\log(|\Delta_K| \mathcal{N}(m)))^2)$ and $N = \lfloor 12n + \log \mathcal{N}(m) + \log |\Delta_K| \rfloor$ is sufficient.

6.4. Probability-density Correspondence

6.4.1. Result

For an ideal set $\mathcal{S} \subseteq \mathcal{I}_K$ consisting of integral ideals, we denote by $\mathcal{S}(t)$ the subset of \mathcal{S} consisting of those integral ideals with norm bounded by $t \in \mathbb{R}_{>0}$, which is made precise in the following lemma. With this notation we will define the *local density* in Definition 6.6.

Definition 6.5. Let \mathcal{S} be an set of integral ideals of \mathcal{O}_K . Then we define $\mathcal{S}(t) := \{\mathfrak{b} \in \mathcal{S} \mid \mathcal{N}(\mathfrak{b}) \leq t\}$ and we define the counting function $|\mathcal{S}(\cdot)| : \mathbb{R}_{>0} \rightarrow \mathbb{N}$ of \mathcal{S} by the following rule:

$$|\mathcal{S}(t)| = |\{\mathfrak{b} \in \mathcal{S} \mid \mathcal{N}(\mathfrak{b}) \leq t\}|.$$

Definition 6.6 (Local density of an ideal set). Let $x > 0$ a positive real number, and let \mathcal{S} be a set of integral ideals of K . We define the local density¹

¹Note that this quantity tends to the ‘natural density’ of the ideal set \mathcal{S} when x goes to infinity, since $|\{\mathfrak{a} \mid \mathcal{N}(\mathfrak{a}) < t\}| \sim \rho \cdot t$ [Ove14, §9.5].

6. Ideal sampling

of \mathcal{S} at x as

$$\delta_{\mathcal{S}}[x] = \min_{t \in [x/e^n, x]} \frac{|\mathcal{S}(t)|}{\rho \cdot t} = \min_{t \in [x/e^n, x]} \frac{|\{\mathfrak{b} \in \mathcal{S} \mid \mathcal{N}(\mathfrak{b}) \leq t\}|}{\rho \cdot t}.$$

Definition 6.7 (Infinity ball). *Let $r > 0$ be a positive number, then we denote*

$$r\mathcal{B}_\infty = \{(x_\sigma)_\sigma \in K_{\mathbb{R}} \mid |x_\sigma| \leq r, \text{ for all } \sigma\}.$$

Definition 6.8. *For a distribution \mathcal{D} on $\text{Div}_{K^{\mathfrak{m}}}^0$, we denote by $[\mathcal{D}] = \mathcal{D}|^{K^{\mathfrak{m}},1}$ the distribution on $\text{Pic}_{K^{\mathfrak{m}}}^0$ obtained by periodizing \mathcal{D} with respect to the subgroup $K^{\mathfrak{m},1} \hookrightarrow \text{Div}_{K^{\mathfrak{m}}}^0$ (see Definition 2.3). In other words,*

$$[\mathcal{D}] = \mathcal{D}|^{K^{\mathfrak{m}},1} = \sum_{\alpha \in K^{\mathfrak{m},1}} \mathcal{D}(\cdot + \langle \alpha \rangle).$$

This distribution $[\mathcal{D}]$ describes exactly the distribution of the Arakelov ray class $[\mathbf{a}]$, where $\mathbf{a} \leftarrow \mathcal{D}$.

The main result of this section shows a close relationship between the *local density* of an ideal set \mathcal{S} and the *probability* that the integral ideal $\beta\mathbf{a}^{-1}$ lies in \mathcal{S} for β sampled appropriately from \mathbf{a} . Here, \mathbf{a} is a Arakelov ray divisor whose Arakelov ray class is uniformly distributed.

Theorem 6.9. *Let $r \geq 8 \cdot n \cdot |\Delta_K|^{\frac{3}{2n}} \cdot \mathcal{N}(\mathfrak{m})^{1/n}$, let $\tau \in (\mathcal{O}_K/\mathfrak{m})^*$ and let $\mathcal{S}^{\mathfrak{m}}$ be a set of integral ideals coprime to \mathfrak{m} with local density $\delta_{\mathcal{S}^{\mathfrak{m}}}[r^n]$ at r^n . Let \mathcal{D} be a distribution on $\text{Div}_{K^{\mathfrak{m}}}^0$ such that $[\mathcal{D}]$ is uniform in $\text{Pic}_{K^{\mathfrak{m}}}^0$. Then*

$$\mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} \left[\Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty} \left[(\alpha) \cdot \text{Exp}(-\mathbf{a}) \in \mathcal{S}^{\mathfrak{m}} \mid \alpha \cdot \text{Exp}(-\mathbf{a}_\infty) \in \tau K^{\mathfrak{m},1} \right] \right] \geq \frac{1}{3} \cdot \delta_{\mathcal{S}^{\mathfrak{m}}}[r^n]. \quad (6.77)$$

where α is uniformly sampled from the finite set $\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty$.

Remark 6.10. *The $1/3$ occurring in Equation (6.77) can be made arbitrarily close to one by increasing the radius $r \in \mathbb{R}$ and slightly increasing the density interval $[x/e^n, x]$ in Definition 6.6. For sake of simplicity we just chose $r \in \mathbb{R}$ and the length of the interval $[x/e^n, x]$ to be minimal to achieve the optimal lower bound up to an explicit constant (i.e., Equation (6.77)).*

Remark 6.11. Theorem 6.9 involves a conditional probability. It is possible, with essentially the same proof technique, to rephrase this theorem in such a way that it concerns the intersection of the events $(\alpha) \cdot \text{Exp}(-\mathbf{a}) \in \mathcal{S}^m$ and $\alpha \cdot \text{Exp}(-\mathbf{a}_\infty) \in \tau K^{m,1}$. The probability then depends as well on the number $\phi(m) = |(\mathcal{O}_K/m)^*| = |K^m/K^{m,1}|$ of elements in $(\mathcal{O}_K/m)^*$. More specifically, for a given $\tau \in (\mathcal{O}_K/m)^*$ one can prove that, under the same conditions as in Theorem 6.9,

$$\begin{aligned} & \Pr_{\substack{\mathbf{a} \leftarrow \mathcal{D} \\ \alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty}} \left[(\alpha) \cdot \text{Exp}(-\mathbf{a}) \in \mathcal{S}^m \text{ and } \alpha \cdot \text{Exp}(-\mathbf{a}_\infty) \in \tau K^{m,1} \right] \\ &= \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} \left[\Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty} \left[(\alpha) \cdot \text{Exp}(-\mathbf{a}) \in \mathcal{S}^m \text{ and } \alpha \cdot \text{Exp}(-\mathbf{a}_\infty) \in \tau K^{m,1} \right] \right] \\ &\geq \frac{1}{3 \cdot \phi(m)} \cdot \delta_{\mathcal{S}^m}[r^n]. \end{aligned} \quad (6.78)$$

6.4.2. Proof Overview of Theorem 6.9

In the following text we prove Theorem 6.9, leaving out details. In the later Section 6.5, which follows the exact same structure as this proof overview, a full proof is given, including all required lemmas.

Simplify the Probability by Fixing a Single Ideal $\mathbf{c} \in \mathcal{S}^m$ and a Single Arakelov Divisor $\mathbf{a} \in \text{Div}_{K^m}^0$

The statement Equation (6.77) of Theorem 6.9 involves two random processes: first the random sampling of $\mathbf{a} \leftarrow \mathcal{D}$, then the uniform sampling of an element $\alpha \in \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty$. It is insightful to focus on the latter random process, that of the element α , for a *fixed* $\mathbf{a} \in \text{Div}_{K^m}^0$.

Also, the probability in Equation (6.77) concerns an entire ideal set \mathcal{S}^m . In this part of the proof, we focus instead on a single ideal $\mathbf{c} \in \mathcal{S}^m$. In other words, we estimate the following probability, for a *fixed* $\mathbf{a} \in \text{Div}_{K^m}^0$ and a *single* integral ideal $\mathbf{c} \in \mathcal{I}_K^m$,

$$\mathbb{P}_{\mathbf{a}, \mathbf{c}} = \Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty} \left[(\alpha) \cdot \text{Exp}(-\mathbf{a}) = \mathbf{c} \mid \alpha \cdot \text{Exp}(-\mathbf{a}_\infty) \in \tau K^{m,1} \right] \quad (6.79)$$

6. Ideal sampling

where the sampling $\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty$ is uniform. In the notation above we leave the dependency on $r \in \mathbb{R}$, $\mathfrak{m} \subseteq \mathcal{O}_K$ and $\tau \in (\mathcal{O}_K/\mathfrak{m})^*$ implicit.

By the law of conditional probability, we have that Equation (6.79) equals

$$\mathbb{P}_{\mathbf{a}, \mathfrak{c}} = \frac{\Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty} \left[\begin{array}{l} (\alpha) \cdot \text{Exp}(-\mathbf{a}) = \mathfrak{c} \\ \text{and} \\ \alpha \cdot \text{Exp}(-\mathbf{a}_\infty) \in \tau K^{\mathfrak{m},1} \end{array} \right]}{\Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty} [\alpha \cdot \text{Exp}(-\mathbf{a}_\infty) \in \tau K^{\mathfrak{m},1}]} \quad (6.80)$$

Focusing on the numerator first, we will prove later, in Lemma 6.12, that

$$\Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty} \left[\begin{array}{l} (\alpha) \cdot \text{Exp}(-\mathbf{a}) = \mathfrak{c} \\ \text{and} \\ \alpha \cdot \text{Exp}(-\mathbf{a}_\infty) \in \tau K^{\mathfrak{m},1} \end{array} \right] = \frac{|\text{Exp}(\mathbf{a} + d(\mathfrak{c}))_\tau^\times \cap r\mathcal{B}_\infty|}{|\text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty|} \quad (6.81)$$

Here, $|\text{Exp}(\mathbf{a} + d(\mathfrak{c}))_\tau^\times \cap r\mathcal{B}_\infty|$ is the number of *generators* of the ideal lattice $\text{Exp}(\mathbf{a})\mathfrak{c}$ that are equivalent to τ modulo \mathfrak{m} (see Definition 6.2) lying in the box $r\mathcal{B}_\infty$. So, essentially, Equation (6.81) counts how many of the $|\text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty|$ elements in the sampling space $r\mathcal{B}_\infty$ actually generate the ideal lattice $\text{Exp}(\mathbf{a})\mathfrak{c}$ and are equivalent to τ modulo \mathfrak{m} .

For the denominator we will prove (see Lemma 6.12) that there exists $\tilde{\tau} \in K_{\mathbb{R}}$ such that

$$\Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty} [\alpha \cdot \text{Exp}(-\mathbf{a}_\infty) \in \tau K^{\mathfrak{m},1}] = \frac{|(\text{Exp}(\mathbf{a})\mathfrak{m} + \tilde{\tau}) \cap r\mathcal{B}_\infty|}{|\text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty|} \quad (6.82)$$

where the sampling $\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty$ is uniform. This equation can be intuitively derived by ignoring the \mathbf{a}_∞ -part. Any element $\alpha \in \text{Exp}(\mathbf{a})$ that satisfies $\alpha \equiv \tau$ modulo \mathfrak{m} must lie in $\text{Exp}(\mathbf{a}) \cap (\mathfrak{m} + \tau)$, which can indeed be rewritten to $\text{Exp}(\mathbf{a})\mathfrak{m} + \tilde{\tau}$ by choosing an $\tilde{\tau} \in \text{Exp}(\mathbf{a})$ that satisfies $\tilde{\tau} \equiv \tau$ modulo \mathfrak{m} .

By combining Equations (6.80) to (6.82) and scratching terms that occur both in the numerator and denominator, one concludes that there exists

$\tilde{\tau} \in \text{Exp}(\mathbf{a})$ such that

$$\begin{aligned} \mathbb{P}_{\mathbf{a}, \mathbf{c}} &= \Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty} \left[(\alpha) \cdot \text{Exp}(-\mathbf{a}) = \mathbf{c} \mid \alpha \cdot \text{Exp}(-\mathbf{a}_\infty) \in \tau K^{\mathbf{m}, 1} \right] \\ &= \frac{|\text{Exp}(\mathbf{a} + d(\mathbf{c}))_\tau^\times \cap r\mathcal{B}_\infty|}{|(\text{Exp}(\mathbf{a})\mathbf{m} + \tilde{\tau}) \cap r\mathcal{B}_\infty|}. \quad (6.83) \end{aligned}$$

Using the estimate

$$|(\text{Exp}(\mathbf{a}) + \tilde{\tau}) \cap r\mathcal{B}_\infty| \approx r^n \cdot 2^{n_{\mathbb{R}}} \cdot (2\pi)^{n_{\mathbb{C}}} / (e^{\deg(\mathbf{a})} \cdot \sqrt{|\Delta_K|})$$

When the radius r is quite large compared to the lattice $\text{Exp}(\mathbf{a}) \subseteq K_{\mathbb{R}}$, one can deduce that for $\mathbf{a} \in \text{Div}_{K^{\mathbf{m}}}$ the number of points in $(\text{Exp}(\mathbf{a}) + \tilde{\tau}) \cap r\mathcal{B}_\infty$ is approximately $\text{Vol}(r\mathcal{B}_\infty) / e^{\deg(\mathbf{a})}$, where $\deg(\cdot)$ is defined in Equation (2.12). More specifically, in Lemma 6.13 we prove that for all $\mathbf{a} \in \text{Div}_{K^{\mathbf{m}}}^0$ and $\tilde{\tau} \in K_{\mathbb{R}}$,

$$|(\text{Exp}(\mathbf{a})\mathbf{m} + \tilde{\tau}) \cap r\mathcal{B}_\infty| \in [e^{-1/4}, e^{1/4}] \cdot r^n \cdot 2^{n_{\mathbb{R}}} \cdot (2\pi)^{n_{\mathbb{C}}} / (\mathcal{N}(\mathbf{m}) \cdot \sqrt{|\Delta_K|}).$$

Applying this to the denominator of Equation (6.83), we directly deduce that

$$\begin{aligned} \mathbb{P}_{\mathbf{a}, \mathbf{c}} &= \Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty} \left[(\alpha) \cdot \text{Exp}(-\mathbf{a}) = \mathbf{c} \mid \alpha \cdot \text{Exp}(-\mathbf{a}_\infty) \in \tau K^{\mathbf{m}, 1} \right] \\ &\in [e^{-1/4}, e^{1/4}] \cdot \frac{\sqrt{|\Delta_K|} \cdot \mathcal{N}(\mathbf{m})}{r^n \cdot 2^{n_{\mathbb{R}}} \cdot (2\pi)^{n_{\mathbb{C}}}} \cdot |\text{Exp}(\mathbf{a} + d(\mathbf{c}))_\tau^\times \cap r\mathcal{B}_\infty| \quad (6.84) \end{aligned}$$

Estimating the probability of sampling a single fixed ideal for a random Arakelov divisor

As already mentioned, Equation (6.77) of Theorem 6.9 involves two random processes, where the first process samples the Arakelov ray divisor $\mathbf{a} \leftarrow \mathcal{D}$ and the second process samples $\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty$ uniformly. Therefore,

6. Ideal sampling

for a fixed integral ideal $\mathfrak{c} \in \mathcal{I}_K^{\mathfrak{m}}$, using Equation (6.84), we have

$$\begin{aligned} & \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} [\mathbb{P}_{\mathbf{a}, \mathfrak{c}}] \\ &= \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} \left[\Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty} \left[(\alpha) \cdot \text{Exp}(-\mathbf{a}) = \mathfrak{c} \mid \alpha \cdot \text{Exp}(-\mathbf{a}_\infty) \in \tau K^{\mathfrak{m}, 1} \right] \right] \\ & \in [e^{-1/4}, e^{1/4}] \cdot \frac{\sqrt{|\Delta_K|} \cdot \mathcal{N}(\mathfrak{m})}{r^n \cdot 2^{n_{\mathbb{R}}} \cdot (2\pi)^{n_{\mathbb{C}}}} \cdot \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} \left[|\text{Exp}(\mathbf{a} + d(\mathfrak{c}))_\tau^\times \cap r\mathcal{B}_\infty| \right]. \quad (6.85) \end{aligned}$$

The number $|\text{Exp}(\mathbf{a} + d(\mathfrak{c}))_\tau^\times \cap r\mathcal{B}_\infty|$ only depends on the Arakelov ray class of $\mathbf{a} \in \text{Div}_{K^{\mathfrak{m}}}^0$

By quite a technical argument (see Lemma 6.14(iii)) one can show that the number of ‘good’ α ’s, $|\text{Exp}(\mathbf{a} + d(\mathfrak{c}))_\tau^\times \cap r\mathcal{B}_\infty|$, only depends on the real number $r \in \mathbb{R}_{>0}$, the Arakelov ray class of the divisor \mathbf{a} and $\tau \in (\mathcal{O}_K/\mathfrak{m})^*$.

Since the distribution \mathcal{D} is assumed to be uniform when projected to the Arakelov ray class group $\text{Pic}_{K^{\mathfrak{m}}}^0$, we can deduce that, for any fundamental domain F of $\text{Pic}_{K^{\mathfrak{m}}}^0$ in $\text{Div}_{K^{\mathfrak{m}}}^0$,

$$\mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} [|\text{Exp}(\mathbf{a} + d(\mathfrak{c}))_\tau^\times \cap r\mathcal{B}_\infty|] = \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{U}(F)} [|\text{Exp}(\mathbf{a} + d(\mathfrak{c}))_\tau^\times \cap r\mathcal{B}_\infty|] \quad (6.86)$$

where $\mathcal{U}(F)$ is the uniform distribution on the fundamental domain F .

By scaling, one can show that $|\text{Exp}(\mathbf{a} + d(\mathfrak{c}))_\tau^\times \cap r\mathcal{B}_\infty| = |\text{Exp}(\mathbf{a} + d^0(\mathfrak{c}))_\tau^\times \cap r\mathcal{N}(\mathfrak{c})^{-1/n}\mathcal{B}_\infty|$. By another technical argument (see Lemma 6.14(i)) one can show that this quantity is non-zero only if $[\mathbf{a} + d^0(\mathfrak{c})] \in [(\tau^{-1})]T^{\mathfrak{m}} \subseteq \text{Pic}_{K^{\mathfrak{m}}}^0$, i.e., if the Arakelov ray class of $\mathbf{a} + d^0(\mathfrak{c})$ lies in a specific coset of the ray unit torus in $\text{Pic}_{K^{\mathfrak{m}}}^0$. We can then deduce that for any fundamental domain $F_{T^{\mathfrak{m}}}$ of $T^{\mathfrak{m}}$ in H ,

$$\begin{aligned} & \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{U}(F)} [|\text{Exp}(\mathbf{a} + d(\mathfrak{c}))_\tau^\times \cap r\mathcal{B}_\infty|] \\ &= \frac{1}{|\text{Cl}_K^{\mathfrak{m}}|} \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{U}(F_{T^{\mathfrak{m}}})} [|\text{Exp}(\mathbf{a})_\tau^\times \cap r\mathcal{N}(\mathfrak{c})^{-1/n}\mathcal{B}_\infty|]. \quad (6.87) \end{aligned}$$

where $\mathcal{U}(F_{T^{\mathfrak{m}}})$ is the uniform distribution on the fundamental domain $F_{T^{\mathfrak{m}}}$.

Taking the logarithmic map into $H = \text{Log } K_{\mathbb{R}}^0$

Applying the logarithmic map on the set $\text{Exp}(\mathbf{a})_{\tau}^{\times} \cap r \cdot \mathcal{N}(\mathfrak{c})^{-1/n} \mathcal{B}_{\infty}$, sends $\text{Exp}(\mathbf{a})_{\tau}^{\times}$ to a shift of the logarithmic ray unit lattice $\Lambda_{K^{\mathfrak{m},1}} = \text{Log}(\mathcal{O}_{K^{\mathfrak{m},1}}^{\times})$ and $r \cdot \mathcal{N}(\mathfrak{c})^{-1/n} \mathcal{B}_{\infty}$ to a simplex $S_{n \log r - \log \mathcal{N}(\mathfrak{c})}$ of volume $C(r, \mathcal{N}(\mathfrak{c}))$, where $S_x = \text{Log}(x \mathcal{B}_{\infty}) \subseteq \text{Log } K_{\mathbb{R}}$ as in Lemma A.1.

The expected value as in Equation (6.87) then equals the average number of points of a randomly shifted logarithmic ray unit lattice into this simplex, which equals $C(r, \mathcal{N}(\mathfrak{c})) / \text{Vol}(T^{\mathfrak{m}})$. The precise value is as follows.

$$\frac{1}{|\text{Cl}_K^{\mathfrak{m}}|} \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{U}(F_{T^{\mathfrak{m}}})} [|\text{Exp}(\mathbf{a})_{\tau}^{\times} \cap r \mathcal{N}(\mathfrak{c})^{-1/n} \mathcal{B}_{\infty}|] = \frac{|\mu_K| \cdot C(r, \mathcal{N}(\mathfrak{c}))}{\phi(\mathfrak{m}) \cdot h_K \cdot R_K} \quad (6.88)$$

Applying the Abel summation formula to get the probability for the ideal set $\mathcal{S}^{\mathfrak{m}}$

By combining Equations (6.85) to (6.88), using the class number formula (see Equation (2.11)) and by the fact that $\frac{\mathcal{N}(\mathfrak{m})}{\phi(\mathfrak{m})} = \frac{|\mathcal{O}_K/\mathfrak{m}|}{|(\mathcal{O}_K/\mathfrak{m})^*|} \geq 1$, one obtains,

$$\begin{aligned} & \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} [\mathbb{P}_{\mathbf{a}, \mathfrak{c}}] \\ &= \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} \left[\Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r \mathcal{B}_{\infty}} \left[(\alpha) \cdot \text{Exp}(-\mathbf{a}) = \mathfrak{c} \mid \alpha \cdot \text{Exp}(-\mathbf{a}_{\infty}) \in \tau K^{\mathfrak{m},1} \right] \right] \\ &\geq e^{-1/4} \cdot \frac{\sqrt{|\Delta_K|} \cdot \mathcal{N}(\mathfrak{m})}{r^n \cdot 2^{n_{\mathbb{R}}} \cdot (2\pi)^{n_{\mathbb{C}}}} \cdot \frac{|\mu_K| \cdot C(r, \mathcal{N}(\mathfrak{c}))}{\phi(\mathfrak{m}) \cdot h_K \cdot R_K} = e^{-1/4} \cdot \frac{C(r, \mathcal{N}(\mathfrak{c}))}{r^n \cdot \rho_K} \cdot \frac{\mathcal{N}(\mathfrak{m})}{\phi(\mathfrak{m})} \\ &\geq e^{-1/4} \cdot \frac{C(r, \mathcal{N}(\mathfrak{c}))}{r^n \cdot \rho_K}, \end{aligned} \quad (6.89)$$

By taking the sum over all $\mathfrak{c} \in \mathcal{S}^{\mathfrak{m}}$, using linearity of the expected value operator, one can achieve the following lower bound.

$$\begin{aligned} & \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} \left[\Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r \mathcal{B}_{\infty}} \left[(\alpha) \cdot \text{Exp}(-\mathbf{a}) \in \mathcal{S}^{\mathfrak{m}} \mid \alpha \cdot \text{Exp}(-\mathbf{a}_{\infty}) \in \tau K^{\mathfrak{m},1} \right] \right] \\ &= \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} \left[\sum_{\mathfrak{c} \in \mathcal{S}^{\mathfrak{m}}} \mathbb{P}_{\mathbf{a}, \mathfrak{c}} \right] = \sum_{\mathfrak{c} \in \mathcal{S}^{\mathfrak{m}}} \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} [\mathbb{P}_{\mathbf{a}, \mathfrak{c}}] \geq e^{-1/4} \sum_{\mathfrak{c} \in \mathcal{S}^{\mathfrak{m}}} \frac{C(r, \mathcal{N}(\mathfrak{c}))}{\rho_K \cdot r^n} \end{aligned} \quad (6.90)$$

6. Ideal sampling

By an application of the Abel summation formula, one can relate the sum $\sum_{\mathfrak{c} \in \mathcal{S}^m} C(r, \mathcal{N}(\mathfrak{c}))$ with an integral involving the counting function $|\mathcal{S}^m(t)| = |\{\mathfrak{c} \in \mathcal{S}^m \mid \mathcal{N}(\mathfrak{a}) \leq t\}|$ of the ideal set \mathcal{S}^m . In fact, we will show that, for some function $f : \mathbb{R} \rightarrow \mathbb{R}$,

$$\sum_{\mathfrak{c} \in \mathcal{S}^m} \frac{C(r, \mathcal{N}(\mathfrak{c}))}{\rho_K \cdot r^n} = \int_{t=1}^{r^n} \frac{|\mathcal{S}^m(t)|}{\rho_K \cdot t} \cdot f(t) dt \geq \delta_{\mathcal{S}^m}[r^n]/2 \quad (6.91)$$

where the last inequality is due to the function $f(t)$ having most of his weight in the interval $[r^n/e^n, r^n]$; precisely the relevant interval for the local density $\delta_{\mathcal{S}^m}[r^n]$ (see Definition 6.6). By combining Equations (6.90) and (6.91), we obtain

$$\mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} \left[\Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty} [(\alpha) \cdot \text{Exp}(-\mathbf{a}) \in \mathcal{S}^m \mid \alpha \cdot \text{Exp}(-\mathbf{a}_\infty) \in \tau K^{m,1}] \right] \geq \delta_{\mathcal{S}^m}[r^n]/3.$$

which finishes the proof.

6.5. Extended Proof of Theorem 6.9

6.5.1. Simplify the Probability by Fixing a Single Ideal $\mathfrak{c} \in \mathcal{S}^m$ and a single Arakelov divisor $\mathbf{a} \in \text{Div}_{K^m}^0$

In this part of the proof we focus on a fixed $\mathbf{a} \leftarrow \text{Div}_{K^m}^0$ in the probabilistic process of Equation (6.77) in Theorem 6.9, leaving the remaining randomness to be the uniform sampling of $\alpha \in \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty$. Furthermore, we concentrate on a fixed ideal $\mathfrak{c} \in \mathcal{S}^m$ as well. By the law of conditional probability, we have

$$\begin{aligned} \mathbb{P}_{\mathbf{a}, \mathfrak{c}} &= \Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty} [(\alpha) \cdot \text{Exp}(-\mathbf{a}) = \mathfrak{c} \mid \alpha \cdot \text{Exp}(-\mathbf{a}_\infty) \in \tau K^{m,1}] \\ &= \frac{\Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty} \left[\begin{array}{c} (\alpha) \cdot \text{Exp}(-\mathbf{a}) = \mathfrak{c} \\ \text{and} \\ \alpha \cdot \text{Exp}(-\mathbf{a}_\infty) \in \tau K^{m,1} \end{array} \right]}{\Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty} [\alpha \cdot \text{Exp}(-\mathbf{a}_\infty) \in \tau K^{m,1}]} \end{aligned} \quad (6.92)$$

In the following lemma we compute the exact values of these probabilities.

Lemma 6.12. *Let $\mathfrak{m} \subseteq \mathcal{O}_K$ be a modulus, let $\tau \in (\mathcal{O}_K/\mathfrak{m})^*$, let $\mathbf{a} \in \text{Div}_{K^{\mathfrak{m}}}^0$ be a fixed Arakelov ray divisor, and let $\mathfrak{c} \in \mathcal{I}_K^{\mathfrak{m}}$ be an integral ideal. Then*

$$\Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_{\infty}} \left[\begin{array}{l} (\alpha) \cdot \text{Exp}(-\mathbf{a}) = \mathfrak{c} \\ \text{and} \\ \alpha \cdot \text{Exp}(-\mathbf{a}_{\infty}) \in \tau K^{\mathfrak{m},1} \end{array} \right] = \frac{|\text{Exp}(\mathbf{a} + d(\mathfrak{c}))_{\tau}^{\times} \cap r\mathcal{B}_{\infty}|}{|\text{Exp}(\mathbf{a}) \cap r\mathcal{B}_{\infty}|}, \quad (6.93)$$

and, there exists some $\tilde{\tau} \in K_{\mathbb{R}}$ such that

$$\Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_{\infty}} \left[\alpha \cdot \text{Exp}(-\mathbf{a}_{\infty}) \in \tau K^{\mathfrak{m},1} \right] = \frac{|(\text{Exp}(\mathbf{a})\mathfrak{m} + \tilde{\tau}) \cap r\mathcal{B}_{\infty}|}{|\text{Exp}(\mathbf{a}) \cap r\mathcal{B}_{\infty}|}, \quad (6.94)$$

where the sampling $\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_{\infty}$ is uniform in both expressions.

Proof. By examining Definition 6.2 closely, noting that $\text{Exp}((\mathbf{a} + d(\mathfrak{c}))_{\mathbf{f}}) = \text{Exp}(\mathbf{a}_{\mathbf{f}}) \cdot \mathfrak{c} \in \mathcal{I}_K^{\mathfrak{m}}$, we see that for all $\alpha \in \text{Exp}(\mathbf{a})$,

$$(\alpha) \cdot \text{Exp}(-\mathbf{a}) = \mathfrak{c} \text{ and } \alpha \cdot \text{Exp}(-\mathbf{a}_{\infty}) \in \tau K^{\mathfrak{m},1} \iff \alpha \in \text{Exp}(\mathbf{a} + d(\mathfrak{c}))_{\tau}^{\times}.$$

As the number of choices for $\alpha \in \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_{\infty}$ equals $|\text{Exp}(\mathbf{a}) \cap r\mathcal{B}_{\infty}|$, the number of good choices equals $|\text{Exp}(\mathbf{a} + d(\mathfrak{c}))_{\tau}^{\times} \cap r\mathcal{B}_{\infty}|$ and since the sampling procedure is uniform, we arrive at the first probability claim. For the second probability claim, write $\mathfrak{a} = \text{Exp}(\mathbf{a}_{\mathbf{f}}) \in \mathcal{I}_K^{\mathfrak{m}}$, for conciseness. We note that for $\alpha \in \text{Exp}(\mathbf{a})$, $\alpha \cdot \text{Exp}(-\mathbf{a}_{\infty}) \in \tau K^{\mathfrak{m},1}$ is equivalent to

$$\alpha \cdot \text{Exp}(-\mathbf{a}_{\infty}) \in \text{Exp}(\mathbf{a}_{\mathbf{f}}) \cap \tau K^{\mathfrak{m},1} = \mathfrak{a} \cap \tau K^{\mathfrak{m},1} = \mathfrak{a}\mathfrak{m} + \tau',$$

where $\tau' \in \mathfrak{a}$ is such that $\tau' \equiv \tau$ modulo \mathfrak{m} (note that \mathfrak{a} and \mathfrak{m} are coprime). This, in turn, is equivalent to

$$\alpha \in \text{Exp}(\mathbf{a}_{\infty})(\mathfrak{a}\mathfrak{m} + \tau') = \text{Exp}(\mathbf{a})\mathfrak{m} + \tilde{\tau}$$

where $\tilde{\tau} = \text{Exp}(\mathbf{a}_{\infty})\tau' \in K_{\mathbb{R}}$, which proves the claim. \square

6.5.2. Estimating the Number of Shifted Lattice Points in a Box

Both Equations (6.93) and (6.94) involve the number of shifted lattice points in the volume $r\mathcal{B}_{\infty}$. For large enough radius r , we can reasonably estimate

6. Ideal sampling

this quantity to be $|(\text{Exp}(\mathbf{a}) + \tilde{\tau}) \cap r\mathcal{B}_\infty| \approx r^n \cdot 2^{n_{\mathbb{R}}} \cdot (2\pi)^{n_{\mathbb{C}}} / (e^{\deg(\mathbf{a})} \cdot \sqrt{|\Delta_K|})$.

Lemma 6.13. *Let $x \geq 1$. For any Arakelov ray divisor $\mathbf{a} \in \text{Div}_{K^m}$, any $r > x \cdot n^2 \cdot |\Delta_K|^{\frac{3}{2n}} \cdot e^{\deg(\mathbf{a})/n}$ and any $\tilde{\tau} \in K_{\mathbb{R}}$, we have*

$$|(\text{Exp}(\mathbf{a}) + \tilde{\tau}) \cap r\mathcal{B}_\infty| \in [e^{-1/x}, e^{1/x}] \cdot \frac{r^n \cdot 2^{n_{\mathbb{R}}} \cdot (2\pi)^{n_{\mathbb{C}}}}{e^{\deg(\mathbf{a})} \cdot \sqrt{|\Delta_K|}},$$

where we note that for $\mathbf{a} \in \text{Div}_{K^m}^0$, i.e., degree-one Arakelov ray divisors, we have $\deg(\mathbf{a}) = 0$.

Proof. Let us write \mathcal{V}_∞ for the Voronoi cell of $\text{Exp}(\mathbf{a})$ around 0 with respect to the infinity norm, i.e., $\mathcal{V}_\infty = \{x \in K_{\mathbb{R}} \mid \|x\|_\infty < \|x - v\|_\infty \text{ for all } v \in \text{Exp}(\mathbf{a})\}$. This is well-known to be a fundamental domain for the lattice $\text{Exp}(\mathbf{a})$ (up to a set of ‘faces’ of measure zero), thus having volume $\det(\text{Exp}(\mathbf{a})) = e^{\deg(\mathbf{a})} \cdot \sqrt{|\Delta_K|}$. Denote $\text{cov}_\infty = \text{cov}_\infty(\text{Exp}(\mathbf{a})) = \max\{\|x\|_\infty \mid x \in \mathcal{V}_\infty\}$ for the covering radius of the lattice $\text{Exp}(\mathbf{a})$ with respect to the infinity norm. Furthermore, denote $\tilde{\tau}_0 \in \mathcal{V}_\infty$ for the unique representative of $\tilde{\tau} + \text{Exp}(\mathbf{a})$ in \mathcal{V}_∞ , implying $\text{Exp}(\mathbf{a}) + \tilde{\tau} = \text{Exp}(\mathbf{a}) + \tilde{\tau}_0$.

The sets $v + \mathcal{V}_\infty$ for $v \in (\text{Exp}(\mathbf{a}) + \tilde{\tau}_0) \cap r\mathcal{B}_\infty$ are disjoint and are included in $K_{\mathbb{R}} \cap (r + 2 \cdot \text{cov}_\infty)\mathcal{B}_\infty$. Hence, by computing the volume of $\cup_{v \in (\text{Exp}(\mathbf{a}) + \tilde{\tau}_0) \cap r\mathcal{B}_\infty} (v + \mathcal{V}_\infty)$ in $K_{\mathbb{R}}$, we obtain

$$\begin{aligned} |(\text{Exp}(\mathbf{a}) + \tilde{\tau}_0) \cap r\mathcal{B}_\infty| \cdot e^{\deg(\mathbf{a})} \cdot \sqrt{|\Delta_K|} &\leq (r + 2 \cdot \text{cov}_\infty)^n \cdot \text{Vol}(K_{\mathbb{R}} \cap \mathcal{B}_\infty) \\ &\leq (r + 2 \cdot \text{cov}_\infty)^n \cdot 2^{n_{\mathbb{R}}} \cdot (2\pi)^{n_{\mathbb{C}}}, \end{aligned}$$

where we used the fact that $\text{Vol}(\mathcal{V}_\infty) = \det(\text{Exp}(\mathbf{a})) = e^{\deg(\mathbf{a})} \cdot \sqrt{|\Delta_K|}$. Observe also that $K_{\mathbb{R}} \cap \mathcal{B}_\infty$ contains some coordinates that are real and other that are complex. Hence, its volume is $2^{n_{\mathbb{R}}} \cdot (2\pi)^{n_{\mathbb{C}}}$ and not 2^n (the 2-dimensional volume of $\{(x, \bar{x}) \in \mathbb{C}^2 \mid |x| \leq 1\}$ is 2π).

In a similar fashion, we can deduce that the set $K_{\mathbb{R}} \cap (r - 2 \cdot \text{cov}_\infty)\mathcal{B}_\infty$ is included in $\cup_{v \in (\text{Exp}(\mathbf{a}) + \tilde{\tau}_0) \cap r\mathcal{B}_\infty} (v + \bar{\mathcal{V}}_\infty)$, where $\bar{\mathcal{V}}_\infty$ is the topological closure

of \mathcal{V}_∞ . The sets $v + \bar{\mathcal{V}}_\infty$ for $v \in (\text{Exp}(\mathbf{a}) + \tilde{\tau}_0) \cap r\mathcal{B}_\infty$ are disjoint up to sets of measure zero, and therefore, by computing volumes, we obtain

$$\begin{aligned} (r - 2 \cdot \text{cov}_\infty)^n \cdot 2^{n\mathbb{R}} \cdot (2\pi)^{n\mathbb{C}} &= (r - 2 \cdot \text{cov}_\infty)^n \cdot \text{Vol}(K_{\mathbb{R}} \cap \mathcal{B}_\infty) \\ &\leq |(\text{Exp}(\mathbf{a}) + \tilde{\tau}_0) \cap r\mathcal{B}_\infty| \cdot e^{\deg(\mathbf{a})} \cdot \sqrt{|\Delta_K|} \end{aligned}$$

Combining the two bounds, one obtains

$$\begin{aligned} \left(1 - \frac{2 \cdot \text{cov}_\infty}{r}\right)^n \cdot \frac{r^n \cdot 2^{n\mathbb{R}} \cdot (2\pi)^{n\mathbb{C}}}{e^{\deg(\mathbf{a})} \cdot \sqrt{|\Delta_K|}} &\leq |(\text{Exp}(\mathbf{a}) + \tilde{\tau}_0) \cap r\mathcal{B}_\infty| \\ &\leq \left(1 + \frac{2 \cdot \text{cov}_\infty}{r}\right)^n \cdot \frac{r^n \cdot 2^{n\mathbb{R}} \cdot (2\pi)^{n\mathbb{C}}}{e^{\deg(\mathbf{a})} \cdot \sqrt{|\Delta_K|}}. \end{aligned}$$

From Lemma 2.22, we know that $\text{cov}_\infty(\text{Exp}(\mathbf{a})) \leq n/2 \cdot |\Delta_K|^{3/2n} \cdot e^{\deg(\mathbf{a})/n}$; so, by assumption, $r \geq 2 \cdot x \cdot n \cdot \text{cov}_\infty$. We obtain the final claim by substituting r and using the inequality $(1 + y/n)^n \leq e^y$ for all $y \in (-1, 1)$. \square

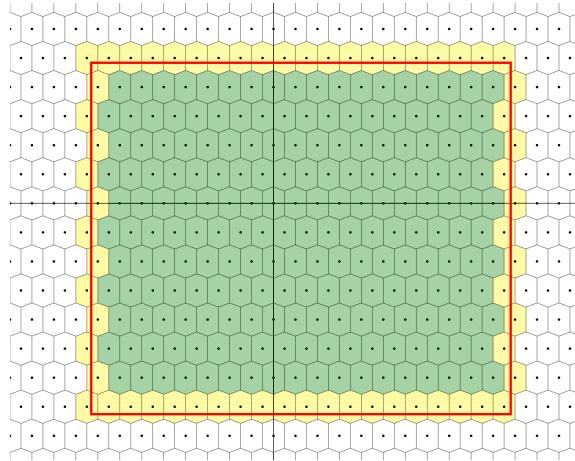


Figure 6.1.: The number of lattice points in the red box is clearly sandwiched by the number of green cells and the number of green and yellow cells together.

Applying above lemma with $x = 4$ and thus $r = 4 \cdot n^2 \cdot |\Delta_K|^{3/2n} \cdot \mathcal{N}(\mathbf{m})^{1/n}$, to Equations (6.93) and (6.94) and substituting them into Equation (6.92),

6. Ideal sampling

one obtains,

$$\begin{aligned}
\mathbb{P}_{\mathbf{a}, \mathfrak{c}} &= \Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty} \left[(\alpha) \cdot \text{Exp}(-\mathbf{a}) = \mathfrak{c} \mid \alpha \cdot \text{Exp}(-\mathbf{a}_\infty) \in \tau K^{\mathfrak{m}, 1} \right] \\
&= \frac{|\text{Exp}(\mathbf{a} + d(\mathfrak{c}))_\tau^\times \cap r\mathcal{B}_\infty|}{|(\text{Exp}(\mathbf{a})\mathfrak{m} + \tilde{\tau}) \cap r\mathcal{B}_\infty|} \\
&\in [e^{-1/4}, e^{1/4}] \cdot \frac{\sqrt{|\Delta_K|} \cdot \mathcal{N}(\mathfrak{m})}{r^n \cdot 2^{n_{\mathbb{R}}} \cdot (2\pi)^{n_{\mathbb{C}}}} \cdot |\text{Exp}(\mathbf{a} + d(\mathfrak{c}))_\tau^\times \cap r\mathcal{B}_\infty| \quad (6.95)
\end{aligned}$$

Here, we use that $\det(\text{Exp}(\mathbf{a})\mathfrak{m}) = e^{\deg(\mathbf{a})} \cdot \sqrt{|\Delta_K|} = \mathcal{N}(\mathfrak{m}) \cdot \sqrt{|\Delta_K|}$, and thus $|(\text{Exp}(\mathbf{a})\mathfrak{m} + \tilde{\tau}) \cap r\mathcal{B}_\infty| \in [e^{-1/4}, e^{1/4}] \cdot \frac{r^n \cdot 2^{n_{\mathbb{R}}} \cdot (2\pi)^{n_{\mathbb{C}}}}{\sqrt{|\Delta_K|} \cdot \mathcal{N}(\mathfrak{m})}$.

6.5.3. Estimating the Probability of Sampling a Single Fixed Ideal for a *Random* Arakelov Divisor

To obtain the probability of sampling a fixed ideal for a *random* Arakelov divisor, one needs to take the weighted sum over the probabilities of sampling a fixed ideal for a fixed Arakelov divisor, where the weights are given by the density \mathcal{D} on $\text{Div}_{K^{\mathfrak{m}}}^0$. In other words, one needs to take the *expected value*. So, for a fixed integral ideal $\mathfrak{c} \in \mathcal{I}_K^{\mathfrak{m}}$, we have

$$\begin{aligned}
&\mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} [\mathbb{P}_{\mathbf{a}, \mathfrak{c}}] \\
&= \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} \left[\Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty} \left[(\alpha) \cdot \text{Exp}(-\mathbf{a}) = \mathfrak{c} \mid \alpha \cdot \text{Exp}(-\mathbf{a}_\infty) \in \tau K^{\mathfrak{m}, 1} \right] \right] \\
&\in [e^{-1/4}, e^{1/4}] \cdot \frac{\sqrt{|\Delta_K|} \cdot \mathcal{N}(\mathfrak{m})}{r^n \cdot 2^{n_{\mathbb{R}}} \cdot (2\pi)^{n_{\mathbb{C}}}} \cdot \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} [|\text{Exp}(\mathbf{a} + d(\mathfrak{c}))_\tau^\times \cap r\mathcal{B}_\infty|]. \quad (6.96)
\end{aligned}$$

where the last approximate equality follows from the linearity of the expectation and Equation (6.95).

6.5.4. The Number $|\text{Exp}(\mathbf{a} + d(\mathfrak{c}))_\tau^\times \cap r\mathcal{B}_\infty|$ Only Depends on the Arakelov Ray Class of $\mathbf{a} \in \text{Div}_{K^{\mathfrak{m}}}^0$

It thus remains to focus on the expected value

$$\mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} [|\text{Exp}(\mathbf{a} + d(\mathfrak{c}))_\tau^\times \cap r\mathcal{B}_\infty|]. \quad (6.97)$$

In the following rather technical lemma we will – among other things – prove that the number of elements in $\text{Exp}(\mathbf{a} + d(\mathbf{c}))_\tau^\times \cap r\mathcal{B}_\infty$ only depends on the Arakelov ray *class* of \mathbf{a} , meaning that we might take the expected value over the uniform distribution over a fundamental domain of $\text{Pic}_{K^m}^0$ in $\text{Div}_{K^m}^0$ in Equation (6.97), as $[\mathcal{D}]$ is uniform over $\text{Pic}_{K^m}^0$, per assumption (see Definition 6.8).

Lemma 6.14. *For all ray divisors $\mathbf{a} \in \text{Div}_{K^m}^0$, elements $\tau, \tau' \in K^m$, ideals $\mathbf{c} \in \mathcal{I}_K^m$ and real numbers $r > 0$ we have the following list of facts.*

- (i) $|\text{Exp}(\mathbf{a})_\tau^\times \cap r\mathcal{B}_\infty| = |\text{Exp}(\mathbf{a} + (\tau'))_{\tau\tau'}^\times \cap r\mathcal{B}_\infty|$, i.e., the number of τ -equivalent ray generators of \mathbf{a} in a fixed box of radius r is equal to the number of $\tau\tau'$ -equivalent ray generators of $\mathbf{a} + (\tau')$ in the same box.
- (ii) $|\text{Exp}(\mathbf{a} + d(\mathbf{c}))_\tau^\times \cap r\mathcal{B}_\infty| = |\text{Exp}(\mathbf{a} + d^0(\mathbf{c}))_\tau^\times \cap \frac{r}{N(\mathbf{c})^{1/n}}\mathcal{B}_\infty|$, so the only difference between the maps d^0 and d is just some appropriate scaling.
- (iii) Writing $\mathbf{a}_\infty = \sum_\nu a_\nu \cdot (\nu) \in H \subseteq \text{Div}_{K^m}^0$, we have

$$|\text{Exp}(\mathbf{a}_\infty)_1^\times \cap r\mathcal{B}_\infty| = |\mu_{K^m, 1}| \cdot |(\Lambda_{K^m, 1} + (a_{\nu_\sigma})_\sigma) \cap S_{\log(r)}|, \quad (6.98)$$

where $S_{\log r} = \{x \in \text{Log } K_{\mathbb{R}} \mid x_\sigma \leq \log(r), \sum_\sigma x_\sigma = 0\}$ is a simplex as in Lemma A.1.

Proof. For part (i), observe that multiplying by $\left(\frac{\sigma(\tau')}{|\sigma(\tau')|}\right)_\sigma \in K_{\mathbb{R}}$ yields a bijection from $\text{Exp}(\mathbf{a})$ to $\text{Exp}(\mathbf{a} + (\tau'))$, preserving the maximum norm. It remains to show that this bijection sends $\text{Exp}(\mathbf{a})_\tau^\times$ to $\text{Exp}(\mathbf{a} + (\tau'))_{\tau'\tau}^\times$. Using Definition 6.2 and assuming $\text{Exp}(\mathbf{a}_f) = \kappa \mathcal{O}_K$ (and therefore $\text{Exp}([\mathbf{a} + (\tau')]_f) = \tau' \kappa \mathcal{O}_K$), we have

$$\begin{aligned} \left(\frac{\sigma(\tau')}{|\sigma(\tau')|}\right)_\sigma \cdot \text{Exp}(\mathbf{a})_\tau^\times &= \left(\frac{1}{|\sigma(\tau')|}\right)_\sigma \cdot (\tau') \cdot \underbrace{\text{Exp}(\mathbf{a}_\infty) \cdot (\kappa \mathcal{O}_K^\times \cap \tau K^{m, 1})}_{\text{Exp}(\mathbf{a})_\tau^\times} \\ &= \underbrace{\left(\frac{1}{|\sigma(\tau')|}\right)_\sigma \cdot \text{Exp}(\mathbf{a}_\infty) \cdot (\tau' \kappa \mathcal{O}_K^\times \cap \tau' \tau K^{m, 1})}_{\text{Exp}((\mathbf{a} + (\tau'))_\infty)} = \text{Exp}(\mathbf{a} + (\tau'))_{\tau'\tau}^\times \end{aligned}$$

6. Ideal sampling

For part (ii), recall that multiplying the ideal lattice $\text{Exp}(d(\mathfrak{c})) = \mathfrak{c} \subseteq K_{\mathbb{R}}$ by the scalar $\mathcal{N}(\mathfrak{c})^{-1/n}$ results in the ideal lattice $\text{Exp}(d^0(\mathfrak{c}))$. Applying this scalar multiplication to the set $\text{Exp}(\mathbf{a} + d(\mathfrak{c})) \cap r\mathcal{B}_{\infty}$ yields a bijective correspondence with $\text{Exp}(\mathbf{a} + d^0(\mathfrak{c})) \cap \frac{r}{\mathcal{N}(\mathfrak{c})^{1/n}}\mathcal{B}_{\infty}$.

In part (iii) it is enough to show that the logarithm $\text{Log} : K_{\mathbb{R}} \rightarrow \text{Log}(K_{\mathbb{R}})$ takes $\text{Exp}(\mathbf{a}_{\infty})_1^{\times}$ to the shifted lattice $\text{Log}(\mathcal{O}_{K^{\mathfrak{m},1}}^{\times}) + (a_{\nu_{\sigma}})_{\sigma} \subset H$ and takes $r\mathcal{B}_{\infty}$ to the simplex $S_{\log(r)} \subset H$. This logarithmic map is $|\mu_{K^{\mathfrak{m},1}}|$ -to-one on $\text{Exp}(\mathbf{a}_{\infty})_1^{\times}$, as it sends roots of unity to the all-one vector in $K_{\mathbb{R}}$, yielding the extra factor $|\mu_{K^{\mathfrak{m},1}}|$ in Equation (6.98). Here, $\mu_{K^{\mathfrak{m},1}} = \mu_K \cap K^{\mathfrak{m},1}$, i.e., the roots of unity in $K^{\mathfrak{m},1}$. \square

As a corollary of Lemma 6.14(i) we deduce that

$$|\text{Exp}(\mathbf{a})_{\tau}^{\times} \cap r\mathcal{B}_{\infty}| = |\text{Exp}(\mathbf{a} + (\kappa))_{\tau}^{\times} \cap r\mathcal{B}_{\infty}|$$

for $\kappa \in K^{\mathfrak{m},1}$, i.e., the number of elements $|\text{Exp}(\mathbf{a})_{\tau}^{\times} \cap r\mathcal{B}_{\infty}|$ only depends on the Arakelov ray class of \mathbf{a} (next to $r \in \mathbb{R}$, \mathfrak{m} and $\tau \in K^{\mathfrak{m}}$). Choose a (measurable) fundamental domain $F \subseteq \text{Div}_{K^{\mathfrak{m}}}^0$ of the quotient group $\text{Pic}_{K^{\mathfrak{m}}}^0$, and put $F_{T^{\mathfrak{m}}} = \{\mathbf{a} \in F \mid [\mathbf{a}] \in T^{\mathfrak{m}}\}$, a fundamental domain of $T^{\mathfrak{m}}$ in $\text{Pic}_{K^{\mathfrak{m}}}^0$. Then, by the assumption that $[\mathcal{D}]$ is uniform on $\text{Pic}_{K^{\mathfrak{m}}}^0$, and writing $\tilde{r} = r\mathcal{N}(\mathfrak{c})^{-1/n}$ we deduce

$$\begin{aligned} & \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} [|\text{Exp}(\mathbf{a} + d(\mathfrak{c}))_{\tau}^{\times} \cap r\mathcal{B}_{\infty}|] = \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} [|\text{Exp}(\mathbf{a} + d^0(\mathfrak{c}))_{\tau}^{\times} \cap \tilde{r}\mathcal{B}_{\infty}|] \\ &= \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{U}(F)} [|\text{Exp}(\mathbf{a} + d^0(\mathfrak{c}))_{\tau}^{\times} \cap \tilde{r}\mathcal{B}_{\infty}|] = \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{U}(F)} [|\text{Exp}(\mathbf{a} + (\tau))_{\tau}^{\times} \cap \tilde{r}\mathcal{B}_{\infty}|] \\ &= \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{U}(F)} [|\text{Exp}(\mathbf{a})_1^{\times} \cap \tilde{r}\mathcal{B}_{\infty}|] = \frac{1}{|\text{Cl}_K^{\mathfrak{m}}|} \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{U}(F_{T^{\mathfrak{m}}})} [|\text{Exp}(\mathbf{a})_1^{\times} \cap \tilde{r}\mathcal{B}_{\infty}|] \end{aligned} \quad (6.99)$$

where the first equality follows from scaling (Lemma 6.14(ii)) and the second one by the fact that the random variable is an Arakelov ray class invariant (Lemma 6.14(i)) and that $[\mathcal{D}]$ is uniform on $\text{Pic}_{K^{\mathfrak{m}}}^0$. The third equality holds because $F + d^0(\mathfrak{c}) - (\tau)$ is a fundamental domain of $\text{Pic}_{K^{\mathfrak{m}}}^0$ in $\text{Div}_{K^{\mathfrak{m}}}^0$ if F is. The fourth equality follows directly from Lemma 6.14(i), and the last equality follows from Definition 6.2. Namely, an Arakelov divisor \mathbf{a} can only have generators if the ideal class of $\text{Exp}(\mathbf{a}_f)$ is trivial, i.e., if $[\mathbf{a}] \in T^{\mathfrak{m}}$. So,

instead, \mathbf{a} can be chosen uniformly from a fundamental domain F_{T^m} of T^m in $\text{Div}_{K^m}^0$, with a correction factor of $\frac{1}{|\text{Cl}_K^m|}$ in the expected value.

6.5.5. Taking the Logarithmic Map into $H = \text{Log } K_{\mathbb{R}}^0$

By taking the Logarithmic image, we find, by Lemma 6.14(iii), that Equation (6.99) equals

$$\frac{|\mu_{K^m,1}|}{|\text{Cl}_K^m|} \cdot \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{U}(F_{T^m})} \left[|\Lambda_{K^m,1} + (a_{\nu_\sigma})_\sigma \cap S_{\log(r) - \log \mathcal{N}(\mathfrak{c})/n}| \right] \quad (6.100)$$

$$= \frac{|\mu_{K^m,1}|}{|\text{Cl}_K^m|} \frac{\text{Vol}(S_{\log(r) - \log \mathcal{N}(\mathfrak{c})/n})}{\text{Vol}(F_{T^m})} \quad (6.101)$$

$$= \frac{|\mu_K| \cdot C(r, \mathcal{N}(\mathfrak{c}))}{|\text{Cl}_K^m| \cdot [\mathcal{O}_K^\times : \mathcal{O}_{K^m,1}^\times] \cdot R_K} = \frac{|\mu_K| \cdot C(r, \mathcal{N}(\mathfrak{c}))}{\phi(\mathfrak{m}) \cdot h_K \cdot R_K}. \quad (6.102)$$

where $C(r, \mathcal{N}(\mathfrak{c})) = \text{Vol}(S_{\log(r) - \log \mathcal{N}(\mathfrak{c})/n}) = (n \log r - \log \mathcal{N}(\mathfrak{c}))^r / r!$ whenever $\mathcal{N}(\mathfrak{c}) \leq r^n$ and zero otherwise; and where $\phi(\mathfrak{m}) = |(\mathcal{O}_K/\mathfrak{m})^*|$. The first inequality of Equation (6.101) follows from Lemma 6.1, the second equality follows from Lemmas A.1 and A.2 and the fact that $\text{Vol}(F_{T^m}) = \text{Vol}(T^m) = [\mathcal{O}_K^\times : \mathcal{O}_{K^m,1}^\times] \cdot |\mu_{K^m,1}| \cdot |\mu_K|^{-1} \cdot \text{Vol}(T)$ (see Lemma 2.16). The third inequality (Equation (6.102)) uses the fact that $|\text{Cl}_K^m| \cdot [\mathcal{O}_K^\times : \mathcal{O}_{K^m,1}^\times] = \phi(\mathfrak{m}) \cdot h_K$ (see Lemma 2.15).

Remark 6.15. Note that all number-theoretic quantities in Equation (6.102) make sense intuitively: one out of h_K random ideal lattices is expected to be principal, the density of units (including roots of unity) is $|\mu_K|/R_K$ and one out of $\phi(\mathfrak{m})$ random elements coprime to \mathfrak{m} equals $\tau \bmod \mathfrak{m}$.

Combining Equations (6.96), (6.99) and (6.100) and the class number formula

6. Ideal sampling

(see Equation (2.11)), we have

$$\begin{aligned}
& \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} [\mathbb{P}_{\mathbf{a}, \mathfrak{c}}] \\
& \quad \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} \left[\Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty} \left[(\alpha) \cdot \text{Exp}(-\mathbf{a}) = \mathfrak{c} \mid \alpha \cdot \text{Exp}(-\mathbf{a}_\infty) \in \tau K^{\mathfrak{m}, 1} \right] \right] \\
& \in [e^{-1/4}, e^{1/4}] \cdot \frac{\sqrt{|\Delta_K|} \cdot \mathcal{N}(\mathfrak{m})}{r^n \cdot 2^{n_{\mathbb{R}}} \cdot (2\pi)^{n_{\mathbb{C}}}} \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} \left[|\text{Exp}(\mathbf{a} + d(\mathfrak{c}))_\tau^\times \cap r\mathcal{B}_\infty| \right] \\
& = [e^{-1/4}, e^{1/4}] \cdot \frac{\sqrt{|\Delta_K|} \cdot \mathcal{N}(\mathfrak{m})}{r^n \cdot 2^{n_{\mathbb{R}}} \cdot (2\pi)^{n_{\mathbb{C}}}} \cdot \frac{|\mu_K| \cdot C(r, \mathcal{N}(\mathfrak{c}))}{\phi(\mathfrak{m}) \cdot h_K \cdot R_K} \\
& = [e^{-1/4}, e^{1/4}] \cdot \frac{C(r, \mathcal{N}(\mathfrak{c}))}{r^n \cdot \rho_K} \cdot \frac{\mathcal{N}(\mathfrak{m})}{\phi(\mathfrak{m})} \geq e^{-1/4} \cdot \frac{C(r, \mathcal{N}(\mathfrak{c}))}{r^n \cdot \rho_K}. \tag{6.103}
\end{aligned}$$

where $C(r, \mathcal{N}(\mathfrak{c})) = (n \log r - \log \mathcal{N}(\mathfrak{c}))^{\mathbb{R}} / \mathbb{R}!$ whenever $\mathcal{N}(\mathfrak{c}) \leq r^n$ and zero otherwise.

6.5.6. Applying the Abel Summation Formula

We have, by Equation (6.103),

$$\begin{aligned}
& \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} \left[\Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty} \left[(\alpha) \text{Exp}(-\mathbf{a}) \in \mathcal{S}^{\mathfrak{m}} \mid \alpha \text{Exp}(\mathbf{a}_\infty) \in \tau K^{\mathfrak{m}, 1} \right] \right] \\
& = \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} \left[\sum_{\mathfrak{c} \in \mathcal{S}^{\mathfrak{m}}} \mathbb{P}_{\mathbf{a}, \mathfrak{c}} \right] = \sum_{\mathfrak{c} \in \mathcal{S}^{\mathfrak{m}}} \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} [\mathbb{P}_{\mathbf{a}, \mathfrak{c}}] \geq \frac{e^{-1/4}}{r^n \cdot \rho_K} \sum_{\substack{\mathfrak{c} \in \mathcal{S}^{\mathfrak{m}} \\ \mathcal{N}(\mathfrak{c}) \leq r^n}} C(r, \mathcal{N}(\mathfrak{c})). \tag{6.104}
\end{aligned}$$

Lemma 6.16. *Let $\mathcal{S}^{\mathfrak{m}} \subseteq \mathcal{I}_K^{\mathfrak{m}}$ a set of integral ideals, let $r > e$, and denote $C(r, \mathcal{N}(\mathfrak{c})) = \frac{(n \log r - \log \mathcal{N}(\mathfrak{c}))^{\mathbb{R}}}{\mathbb{R}!}$. Then*

$$\frac{1}{r^n \cdot \rho_K} \sum_{\substack{\mathfrak{c} \in \mathcal{S}^{\mathfrak{m}} \\ \mathcal{N}(\mathfrak{c}) \leq r^n}} C(r, \mathcal{N}(\mathfrak{c})) \geq \frac{1}{2} \cdot \delta_{\mathcal{S}^{\mathfrak{m}}} [r^n]$$

Proof. We apply the Abel partial summation formula with $a_{N, \mathcal{S}^{\mathfrak{m}}} := |\{\mathfrak{c} \in$

$\mathcal{S}^m \mid \mathcal{N}(\mathfrak{c}) = N\}$ and $C(r, N) := \frac{(n \log r - \log N)^{\mathfrak{r}}}{\mathfrak{r}!}$, whose derivative equals

$$\begin{aligned} \frac{d}{dN} C(r, N) \Big|_t &= -\frac{(n \log r - \log t)^{\mathfrak{r}-1}}{t \cdot (\mathfrak{r}-1)!} \\ &= \frac{-r^n}{t \cdot (\mathfrak{r}-1)!} \cdot \left[\frac{d}{dN} \Gamma(\mathfrak{r}, n \log r - \log N) \right]_t, \end{aligned}$$

where $\Gamma(\mathfrak{r}, x) = \int_x^\infty u^{\mathfrak{r}-1} e^{-u} du$ is the upper incomplete Gamma function. Recall that $|\mathcal{S}^m(t)| = \sum_{N \leq t} a_{N, \mathcal{S}^m}$. A typical application of the Abel summation formula yields

$$\begin{aligned} &\frac{1}{r^n \cdot \rho_K} \sum_{\substack{\mathfrak{c} \in \mathcal{S}^m \\ \mathcal{N}(\mathfrak{c}) \leq r^n}} C(r, \mathcal{N}(\mathfrak{c})) \\ &= \frac{1}{r^n \cdot \rho_K} \sum_{1 \leq N \leq r^n} a_{N, \mathcal{S}^m} \cdot C(r, N) \\ &= - \int_{t=1}^{r^n} \frac{|\mathcal{S}^m(t)|}{\rho_K \cdot r^n} \cdot \left[\frac{d}{dN} C(r, N) \Big|_{N=t} \right] dt \\ &= \frac{1}{(\mathfrak{r}-1)!} \int_{t=1}^{r^n} \frac{|\mathcal{S}^m(t)|}{\rho_K \cdot t} \cdot \left[\frac{d}{dN} \Gamma(\mathfrak{r}, n \log r - \log N) \Big|_{N=t} \right] dt, \end{aligned} \quad (6.105)$$

Using Definition 6.6 about ideal density and the fact that the integrand is positive, Equation (6.105) is lower bounded by

$$\begin{aligned} &\frac{1}{(\mathfrak{r}-1)!} \int_{t=(r/e)^n}^{r^n} \frac{|\mathcal{S}^m(t)|}{\rho_K \cdot t} \cdot \left[\frac{d}{dN} \Gamma(\mathfrak{r}, n \log r - \log N) \Big|_{N=t} \right] dt \\ &\geq \frac{\delta_{\mathcal{S}^m}[r^n]}{(\mathfrak{r}-1)!} \int_{t=(r/e)^n}^{r^n} \left[\frac{d}{dN} \Gamma(\mathfrak{r}, n \log r - \log N) \Big|_{N=t} \right] dt \geq \frac{1}{2} \cdot \delta_{\mathcal{S}^m}[r^n], \end{aligned} \quad (6.106)$$

where the last inequality (Equation (6.106)) follows from the definition of the upper incomplete Gamma function,

$$\begin{aligned} &\frac{1}{(\mathfrak{r}-1)!} \int_{t=(r/e)^n}^{r^n} \left(\frac{d}{dt} \Gamma(\mathfrak{r}, n \log r - \log N) \Big|_{N=t} \right) dt \\ &= \frac{1}{(\mathfrak{r}-1)!} \cdot (\Gamma(\mathfrak{r}, 0) - \Gamma(\mathfrak{r}, n)) = 1 - e^{-n} \sum_{k=0}^{\mathfrak{r}-1} \frac{n^k}{k!} \geq 1/2, \end{aligned}$$

where we used the fact that $e^{-n} \sum_{k=0}^{\mathfrak{r}-1} \frac{n^k}{k!}$ equals the probability that a Poisson distribution with parameter n yields at most $\mathfrak{r}-1 \leq n-1$ occurrences, which is bounded by a half. \square

6. Ideal sampling

We conclude that

$$\begin{aligned}
& \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} [\mathbb{P}_{\mathbf{a}, \mathbf{c}}] \\
&= \mathbb{E}_{\mathbf{a} \leftarrow \mathcal{D}} \left[\Pr_{\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r \mathcal{B}_\infty} \left[(\alpha) \text{Exp}(-\mathbf{a}) \in \mathcal{S}^m \mid \alpha \text{Exp}(-\mathbf{a}_\infty) \in \tau K^{m,1} \right] \right] \\
&\geq \frac{e^{-1/4}}{r^n \cdot \rho_K} \sum_{\substack{\mathbf{c} \in \mathcal{S}^m \\ \mathcal{N}(\mathbf{c}) \leq r^n}} C(r, \mathcal{N}(\mathbf{c})) \geq \delta_{\mathcal{S}^m}[r^n]/3.
\end{aligned}$$

This concludes the proof of Theorem 6.9.

Remark 6.17. As already mentioned, the fraction $\frac{1}{3}$ before $\delta_{\mathcal{S}^m}[r^n]$ can be made arbitrarily close to 1. In order to achieve that, we need to enlarge the ‘ideal density’ interval in Definition 6.6 to $[x/e^{2n}, x]$ and we need to increase the radius $r \in \mathbb{R}_{>0}$ in Lemma 6.13.

In the case of this larger density interval, the Poisson distribution in above proof changes into a Poisson distribution with parameter $2n$, but with the same bound ($\mathbf{r} - 1 \leq n - 1$) on the occurrences. This yields an exponential instead of a constant bound. Increasing the radius r by an exponential factor 2^n also yields an exponential bound on the error. So, by implementing these changes, one can obtain a lower bound on the probability in Theorem 6.9 of $(1 - O(e^{-n})) \cdot \delta_{\mathcal{S}^m}[r^n]$, which is exponentially close to optimal.

6.6. Ideal Sampling

6.6.1. Sampling in a Box

In this section, we explain how one can efficiently sample in a (distorted) infinity box, provided that all the dimensions of the box are sufficiently large. More precisely, let $(r_\sigma)_\sigma \in K_{\mathbb{R}}$ be such that $r_\sigma > 0$ for all coordinates. We let $(r_\sigma)_\sigma \mathcal{B}_\infty$ denote the distorted box

$$(r_\sigma)_\sigma \mathcal{B}_\infty := \{(x_\sigma)_\sigma \in K_{\mathbb{R}} \mid |x_\sigma| \leq r_\sigma, \forall \sigma\}.$$

Proposition 6.18. *Let $\mathfrak{a} \subset \mathcal{O}_K$ be an ideal represented by a basis $M_{\mathfrak{a}}$, let $\beta \in \mathcal{O}_K$ be a shift and let $(r_{\sigma})_{\sigma} \in K_{\mathbb{R}}$ be such that $r_{\sigma} > 0$ for all σ . Assume that for all embeddings σ , it holds that $r_{\sigma} \geq 4 \cdot 2^n \cdot n^{3/2} \cdot |\Delta_K|^{3/(2n)} \cdot \mathcal{N}(\mathfrak{a})^{1/n}$. Then, there exists an algorithm sampling uniformly in $(\mathfrak{a} + \beta) \cap (r_{\sigma})_{\sigma} \mathcal{B}_{\infty}$ using time $O(n^6 \log(|M_{\mathfrak{a}}|)^3)$, where $|M_{\mathfrak{a}}|$ denotes the length of the longest basis vector of $M_{\mathfrak{a}}$.*

Remark 6.19. *This lemma can be seen as the ‘algorithmization’ of the ideas in the very similar Lemma 6.13. In that particular lemma (see also Figure 6.1), an estimation is made of the number of lattice points in a box, where Voronoi cells are used as the fundamental domain.*

In the proof of this lemma, we sample a random element in the ambient vector space of the lattice that also lies in the predescribed box $(r_{\sigma})_{\sigma} \mathcal{B}_{\infty}$. Then, we use a ‘rounding algorithm’ to round this real vector to an actual lattice point in \mathfrak{a} . Such a rounding algorithm needs a fundamental domain of the lattice \mathfrak{a} , which can be computed by means of the LLL-algorithm. This might yield quite a skewed fundamental domain, hence the slightly worse requirements on the parameters, compared to Lemma 6.13.

Proof. The algorithm first computes (in polynomial time) an LLL reduced basis (a_1, \dots, a_n) of \mathfrak{a} from $M_{\mathfrak{a}}$. This basis satisfies $\|a_i\| \leq 2^n \lambda_n(\mathfrak{a}) \leq 2^n \cdot \sqrt{n} \cdot |\Delta_K|^{3/(2n)} \cdot \mathcal{N}(\mathfrak{a})^{1/n}$ (using Lemma 2.22). After that, reduce $\beta \in \mathcal{O}_K$ modulo this LLL reduced basis (a_1, \dots, a_n) of \mathfrak{a} , yielding $\tilde{\beta} \in \beta + \mathfrak{a}$. That is, write $\beta = \sum_i t_i a_i$ and put $\tilde{\beta} = \sum_i (t_i - \lfloor t_i \rfloor) a_i$.

Denoting $D := \sum_i \|a_i\|_{\infty} \leq 2^n \cdot \sqrt{n} \cdot |\Delta_K|^{3/(2n)} \cdot \mathcal{N}(\mathfrak{a})^{1/n}$ for the sum of the infinity norms of the basis vectors a_i , we have $\|\tilde{\beta}\|_{\infty} \leq D$. Also, by assumption on $(r_{\sigma})_{\sigma}$, it holds that $r_{\sigma} \geq 4nD$ for every embedding σ .

To sample a uniform element $v \in \mathfrak{a} \cap (r_{\sigma})_{\sigma} \mathcal{B}_{\infty}$, the algorithm goes through the following steps. It first samples a uniform element $t = (t_{\sigma})_{\sigma} \in (r_{\sigma} + 2D)_{\sigma} \mathcal{B}_{\infty}$. This can be done in time $\text{poly}(n, \log(\max_{\sigma} r_{\sigma} + 2D))$, by sampling every first $r_{\mathbb{R}} + r_{\mathbb{C}}$ coordinates of $t \in K_{\mathbb{R}}$ independently, and defining the last $r_{\mathbb{C}}$ ones appropriately in order to have $t \in K_{\mathbb{R}}$. The algorithm then writes $t = \sum_i t_i a_i$ with $t_i \in \mathbb{R}$ (the vector t lies in the real span of \mathfrak{a}) and puts

6. Ideal sampling

$v = \sum_i \lfloor t_i \rfloor a_i + \tilde{\beta}$, which lies in $\beta + \mathfrak{a}$. Finally, the algorithm outputs v if $v \in (r_\sigma)_\sigma \mathcal{B}_\infty$, otherwise it restarts.

Let us first show that the distribution of v is indeed uniform in $(\mathfrak{a} + \beta) \cap (r_\sigma)_\sigma \mathcal{B}_\infty$. Let us define $\mathcal{P} = \{\sum_i x_i a_i, x_i \in [-1/2, 1/2]\}$ the fundamental parallelepiped associated to the basis (a_1, \dots, a_n) . It holds that for all $x \in \mathcal{P}$, we have $\|x\|_\infty \leq D$.

The probability of sampling $v = \alpha + \tilde{\beta} \in (r_\sigma)_\sigma \mathcal{B}_\infty$ for $\alpha \in \mathfrak{a}$ via the above procedure is equal to the probability of sampling $t \in \alpha + \mathcal{P} \subseteq (r_\sigma + 2D)_\sigma \mathcal{B}_\infty$. This probability is equal to $\text{Vol}(\mathcal{P}) / \text{Vol}((r_\sigma + 2D)_\sigma \mathcal{B}_\infty \cap K_{\mathbb{R}})$, which does not depend on $\alpha \in \mathfrak{a}$. We conclude that above sampling procedure yields $v = \alpha + \tilde{\beta}$ that are uniformly distributed in $(\mathfrak{a} + \beta) \cap (r_\sigma)_\sigma \mathcal{B}_\infty$. The running time of the algorithm is dominated by the LLL-reduction of $M_{\mathfrak{a}}$, which takes time $O(n^6 \log(|M_{\mathfrak{a}}|)^3)$, where $|M_{\mathfrak{a}}|$ denotes the length of the longest basis vector of $M_{\mathfrak{a}}$.

To conclude the proof, we show that the success probability of the algorithm is constant. Indeed, observe that whenever $t = \sum_i t_i a_i \in (r_\sigma - 2D)_\sigma \mathcal{B}_\infty$, then we have $v = \sum_i \lfloor t_i \rfloor a_i + \tilde{\beta} \in (r_\sigma)_\sigma \mathcal{B}_\infty$ (since $\|t - v\|_\infty \leq D$ and $\|\tilde{\beta}\| \leq D$), and so the algorithm succeeds. The success probability of the algorithm is then at least

$$\frac{\text{Vol}((r_\sigma - 2D)_\sigma \mathcal{B}_\infty \cap K_{\mathbb{R}})}{\text{Vol}((r_\sigma + 2D)_\sigma \mathcal{B}_\infty \cap K_{\mathbb{R}})} = \prod_{\sigma} \left(\frac{1 - 2D/r_\sigma}{1 + 2D/r_\sigma} \right)^n \geq \left(\frac{1 - \frac{1}{2n}}{1 + \frac{1}{2n}} \right)^n \geq 1/3,$$

where we used the fact that $2nD/r_\sigma \leq 1/(2n)$ for any σ . This concludes the proof. \square

6.6.2. The Sampling Algorithm

Definition 6.20. We denote by \mathcal{S}_B the set of B -smooth ideals, i.e.,

$$\mathcal{S}_B = \{\mathfrak{a} \text{ ideal of } \mathcal{O}_K \mid \text{for any prime ideal } \mathfrak{p} \mid \mathfrak{a} \text{ holds } \mathcal{N}(\mathfrak{p}) \leq B\}$$

Algorithm 7: Sampling of $\beta \in \mathfrak{b}$ such that $\beta \equiv \tau$ modulo \mathfrak{m}

Require:

- A modulus $\mathfrak{m} \subseteq \mathcal{O}_K$,
- An ideal $\mathfrak{b} \subseteq \mathcal{O}_K$ coprime with \mathfrak{m} ,
- An element $\tau \in \mathcal{O}_K$ coprime with \mathfrak{m} ,
- An error parameter $\varepsilon > 0$.

Ensure: An element $\beta \in \mathfrak{b}$ such that

- $\beta \equiv \tau$ modulo \mathfrak{m} ,
- $|\mathcal{N}(\beta)| \leq \mathcal{N}(\mathfrak{b}) \cdot B^N \cdot r^n$, where $r = 4 \cdot 2^n \cdot n^{3/2} \cdot |\Delta_K|^{\frac{3}{2n}} \cdot \mathcal{N}(\mathfrak{m})^{\frac{1}{n}}$, where $B = \tilde{O}\left(n^4[\log \log(1/\varepsilon)]^2 + n^2[\log(|\Delta_K| \mathcal{N}(\mathfrak{m}))]^2\right)$ and $N = \lfloor 8n + \log \mathcal{N}(\mathfrak{m}) + \log |\Delta_K| + 2 \log(1/\varepsilon) \rfloor$ as in Theorem 6.3.

- 1: Multiply \mathfrak{b} by N random prime ideals coprime with \mathfrak{m} and that have a norm bounded by B , obtaining $\tilde{\mathfrak{b}} = \mathfrak{b} \cdot \prod_j \mathfrak{p}_j$.
- 2: Sample a Gaussian distortion $(x_\sigma)_\sigma \in H \subseteq \log K_{\mathbb{R}}$ with parameter $s = 1/n^2$ and define the $(e^{x_\sigma})_\sigma$ -distorted box $\tilde{\mathcal{B}} = (e^{x_\sigma} \cdot r \cdot \mathcal{N}(\tilde{\mathfrak{b}})^{1/n})_\sigma \mathcal{B}_\infty$.
- 3: Compute $\tilde{\tau} \in \tilde{\mathfrak{b}}$ such that $\tilde{\tau} \equiv \tau$ modulo \mathfrak{m} .
- 4: Sample an element $\beta \in (\tilde{\mathfrak{b}}\mathfrak{m} + \tilde{\tau}) \cap \tilde{\mathcal{B}} = \tilde{\mathfrak{b}} \cap (\mathfrak{m} + \tau) \cap \tilde{\mathcal{B}}$ uniformly random following the algorithm from Proposition 6.18.
- 5: **return** β .

Theorem 6.21 (ERH). *Let \mathcal{S} be any set of integral ideals, let $\mathfrak{m} \subseteq \mathcal{O}_K$ be any ideal modulus, let $\mathfrak{b} \subseteq \mathcal{O}_K$ be an integral ideal coprime with \mathfrak{m} and let $\tau \in (\mathcal{O}_K/\mathfrak{m})^*$ be any invertible element modulo \mathfrak{m} . Let, furthermore, $r \geq 4 \cdot 2^n \cdot n^{3/2} \cdot |\Delta_K|^{\frac{3}{2n}} \cdot \mathcal{N}(\mathfrak{m})^{\frac{1}{n}}$ and let $\varepsilon > 0$ be an error parameter.*

Then, assuming the Extended Riemann Hypothesis, Algorithm 7 outputs in time $T = \text{poly}(\log |\Delta_K|, \text{size}(\mathcal{N}(\mathfrak{b})), \log(1/\varepsilon), \log(\mathcal{N}(\mathfrak{m})))$ an element $\beta \in \mathfrak{b}$ such that

- $(\beta)/\mathfrak{b} \in (\mathcal{S} \cdot \mathcal{S}_B) \cap \mathcal{I}_K^{\mathfrak{m}}$,
- $\beta \equiv \tau \pmod{\mathfrak{m}}$

6. Ideal sampling

with probability at least $\frac{1}{3}\delta_{\mathcal{S}}[r^n] - \varepsilon$.

Here, $B = \tilde{O}\left(n^2 \cdot [n^2 \cdot (\log \log(1/\varepsilon))^2 + (\log(|\Delta_K| \mathcal{N}(\mathfrak{m})))^2]\right)$.

Proof. We split the proof into two parts. We start with the proof of correctness and the success probability and finish with the proof of the polynomial running time.

(*Correctness and success probability*). By Lemma 6.22, which we will treat later, the ideal-element pair $((\beta)/\tilde{\mathfrak{b}}, \beta) \in \mathcal{I}_K^{\mathfrak{m}} \times \mathcal{O}_K$ from Algorithm 7 is distributed as $((\alpha) \text{Exp}(-\mathbf{a}), \alpha \text{Exp}(-\mathbf{a}_\infty))$ with $\mathbf{a} \leftarrow \mathcal{W} = \mathcal{W}(N, B, s) + d^0(\mathfrak{b})$ and $\alpha \leftarrow \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_\infty$ uniformly. Here $\mathcal{W} = \mathcal{W}(N, B, s) + d^0(\mathfrak{b})$ is the random walk distribution starting on the point $d^0(\mathfrak{b}) \in \text{Div}_{K^{\mathfrak{m}}}^0$ (see Definition 4.1).

For the random walk distribution $\mathcal{W} = \mathcal{W}(N, B, s) + d^0(\mathfrak{b})$ on $\text{Div}_{K^{\mathfrak{m}}}^0$ with these parameters holds that $[\mathcal{W}]$ on $\text{Pic}_{K^{\mathfrak{m}}}^0$ is ε -close to uniform in the total variation distance. So, allowing an error of ε we may as well assume that \mathbf{a} instead comes from a distribution \mathcal{D} on $\text{Div}_{K^{\mathfrak{m}}}^0$ that satisfies $[\mathcal{D}] = \mathcal{U}(\text{Pic}_{K^{\mathfrak{m}}}^0)$ (see Lemma 6.23).

By applying Theorem 6.9, one then obtains that the expected probability (over the randomness of $\mathbf{a} \in \text{Div}_{K^{\mathfrak{m}}}^0$) that $(\beta)/\tilde{\mathfrak{b}} = (\alpha) \text{Exp}(-\mathbf{a}) \in \mathcal{S} \cap \mathcal{I}_K^{\mathfrak{m}}$ given that $\beta = \alpha \text{Exp}(-\mathbf{a}_\infty) \equiv \tau \pmod{\mathfrak{m}}$ is at least $\frac{1}{3}\delta_{\mathcal{S}}[r^n] - 2^{-n}$. From the fact that $\tilde{\mathfrak{b}} = \mathfrak{b} \cdot \prod_j \mathfrak{p}_j$ with $\mathfrak{p}_j \nmid \mathfrak{m}$ and $\mathcal{N}(\mathfrak{p}_j) \leq B$, we have that $(\beta)/\mathfrak{b} \in (\mathcal{S} \cdot \mathcal{S}_B) \cap \mathcal{I}_K^{\mathfrak{m}}$ in that case, and the result follows.

(*Running time*). Note that $\log B$ and N are $\text{poly}(\log |\Delta_K|, \text{size}(\mathcal{N}(\mathfrak{b})), \log(1/\varepsilon), \log(\mathcal{N}(\mathfrak{m})))$, so any complexity polynomially involving $\log B$ and N must be polynomial in the size of the input as well. In the following complexity analysis, any complexity that is within $\text{poly}(\log |\Delta_K|, \text{size}(\mathcal{N}(\mathfrak{b})), \log(1/\varepsilon), \log(\mathcal{N}(\mathfrak{m})))$ we will call ‘polynomial in the size of the input’.

For the running time, we go through steps 1 to 4 of Algorithm 7. Step 1 involves the sampling of N primes, which, by Lemma 2.14 takes $O(N \cdot n^2 \log^2 B)$ time, clearly polynomial in the size of the input; the fact that

the primes need to be coprime with \mathfrak{m} does not give a significant overhead². Multiplication of two ideals can be done by LLL-reducing the $n^2 \times n$ matrix involving all products of the \mathbb{Z} -generators of the respective ideals, taking time at most $\tilde{O}(n^{8+\eta} \log(M)^{1+\eta})$ [NS16] for any $\eta > 0$, where M is the maximum entry of the matrix involved³. This multiplication is done with N ideals, which means that the total time of this ideal multiplication is polynomial in the size of the input. An alternative way to see this is by using the two-element representation of ideals (e.g., [CS08, §4.7]).

Step 2 requires to sample a Gaussian in $H = \text{Log } K_{\mathbb{R}}^0$, which can be done by inverse transform sampling, without a significant running time. The estimation of the time required for sampling in the box $(e^{x\sigma} \cdot r \cdot \mathcal{N}(\tilde{\mathfrak{b}})^{1/n})_{\sigma} \mathcal{B}_{\infty}$ is deferred to step 4.

In step 3 one only needs to compute $\beta \in \tilde{\mathfrak{b}}$ and $\mu \in \mathfrak{m}$ such that $\beta + \mu = 1$. In that case $\tilde{\tau} = \beta\tau$ suffices. Such a pair $(\beta, \mu) \in \tilde{\mathfrak{b}} \times \mathfrak{m}$ can be found by applying the Hermite normal form to the concatenated basis matrices of $\tilde{\mathfrak{b}}$ and \mathfrak{m} [Coh99, Prop. 1.3.1]. This requires $\tilde{O}(n^5 \log(M)^2)$ time [SL96], where M is the maximum entry occurring in the basis matrices.

Step 4 requires the sampling-in-a-box algorithm described in Proposition 6.18 which requires $O(n^6 \log(|\Delta_K| \mathcal{N}(\mathfrak{m}\tilde{\mathfrak{b}}))^3) = O(n^6 \log(|\Delta_K| \mathcal{N}(\mathfrak{m}) B^N)^3)$ time.

Clearly all steps require time at most polynomial in the size of the input, which proves the time complexity claim. \square

Above proof needs the results of Lemma 6.22 and Lemma 6.23. The first proves the fact that Algorithm 7 mimics a random walk, and the second shows that the random walk distribution on $\text{Div}_{K^{\mathfrak{m}}}^0$ is close to a distribution \mathcal{D} for which $[\mathcal{D}]$ is uniform on $\text{Pic}_{K^{\mathfrak{m}}}^0$. After these two lemmas, the proof is completed.

²In the sampling procedure in Lemma 2.14, the first step is sampling a random integer p in $[0, B]$. In this particular step one can avoid primes dividing \mathfrak{m} by simply compute the greatest common divisor of p and $\mathcal{N}(\mathfrak{m})$. This only gives a non-significant overhead compared to the full algorithm in Lemma 2.14.

³This time estimate is from Neumaier and Stehlé [NS16], instantiated with $\beta = \log \max_i \|\mathbf{b}_i\| \leq \log(nM)$ and lattice dimension n^2 .

Lemma 6.22 (Algorithm 7 mimicks a random walk). *Let $\mathfrak{m} \subseteq \mathcal{O}_K$ a modulus, let N, B, s and r as in Algorithm 7 and let $\mathcal{W} = \mathcal{W}(N, B, s)$ be the random walk distribution on $\text{Div}_{K^{\mathfrak{m}}}^0$ (see Definition 4.1). Let \mathcal{W}_r be the distribution on $K_{\mathbb{R}} \times \text{Div}_{K^{\mathfrak{m}}}^0$ obtained by sampling $\mathbf{a} \leftarrow \mathcal{W}(N, B, s) + d^0(\mathfrak{b})$ and subsequently sampling $\alpha \in \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_{\infty}$ uniformly.*

Then the pair $((\beta)/\tilde{\mathfrak{b}}, \beta) \in \mathcal{I}_K^{\mathfrak{m}} \times \mathcal{O}_K$ obtained by running Algorithm 7 follows the exact same distribution as $((\alpha) \text{Exp}(-\mathbf{a}), \alpha \text{Exp}(-\mathbf{a}_{\infty}))$ with $(\alpha, \mathbf{a}) \leftarrow \mathcal{W}_r$.

Proof. The difference in the sampling procedure consists of where the disturbance of the ‘infinite places’ happens. In the case of the random walk, the disturbance happens on the divisor, whereas in Algorithm 7 the disturbance happens on the box to be sampled in. We will show that this does not matter for the end distribution.

Both the distribution \mathcal{W} and Algorithm 7 involve the following two random processes: picking N uniformly random primes from

$$\{\mathfrak{p} \in \mathcal{I}_K^{\mathfrak{m}} \text{ prime} \mid \mathcal{N}(\mathfrak{p}) \leq B\}$$

and sampling a Gaussian $(x_{\sigma})_{\sigma} \in H$; both with the exact same parameters. Without loss of generality, we can therefore focus on one fixed sample $\{\mathfrak{p}_j \mid 1 \leq j \leq N\}$ of primes and one fixed vector $(x_{\sigma})_{\sigma} \in H$.

This means that we consider the fixed $\mathbf{a} = \sum_{j=1}^N (\mathfrak{p}_j) + \sum_{\nu} x_{\sigma_{\nu}} (\nu) + d^0(\mathfrak{b}) \in \text{Div}_{K^{\mathfrak{m}}}^0$ for the procedure involving \mathcal{W}_r and the fixed ideal $\tilde{\mathfrak{b}} = \mathfrak{b} \prod_{j=1}^N \mathfrak{p}_j$ and distortion $(e^{-x_{\sigma}})_{\sigma}$ for the procedure involving Algorithm 7. Then, writing $\tilde{b} = \mathcal{N}(\tilde{\mathfrak{b}})^{1/n}$,

$$\text{Exp}(\mathbf{a}) = (e^{x_{\sigma}})_{\sigma} \tilde{\mathfrak{b}} / \tilde{b}, \quad \text{Exp}(\mathbf{a}_f) = \tilde{\mathfrak{b}} \quad \text{and} \quad \text{Exp}(\mathbf{a}_{\infty}) = (e^{x_{\sigma}})_{\sigma} / \tilde{b}$$

Thus, $\alpha \text{Exp}(-\mathbf{a}_{\infty})$ for uniformly random $\alpha \in \text{Exp}(\mathbf{a}) \cap r\mathcal{B}_{\infty}$ is distributed as

$$\begin{aligned} \text{Exp}(-\mathbf{a}_{\infty}) \cdot \mathcal{U}(\text{Exp}(\mathbf{a}) \cap r\mathcal{B}_{\infty}) &= (e^{-x_{\sigma}})_{\sigma} \cdot \tilde{b} \cdot \mathcal{U}((e^{x_{\sigma}})_{\sigma} \tilde{\mathfrak{b}} / \tilde{b} \cap r\mathcal{B}_{\infty}) \\ &= \mathcal{U}(\tilde{\mathfrak{b}} \cap (e^{-x_{\sigma}})_{\sigma} \cdot \tilde{b} \cdot r\mathcal{B}_{\infty}) \end{aligned}$$

which is exactly the distribution of $\beta \in \tilde{\mathfrak{b}}$ in Algorithm 7 for fixed $\tilde{\mathfrak{b}}$. It follows that $(\alpha) \text{Exp}(-\mathbf{a}) = (\alpha) \text{Exp}(-\mathbf{a}_\infty) \text{Exp}(-\mathbf{a}_f) = (\alpha) \text{Exp}(-\mathbf{a}_\infty)/\tilde{\mathfrak{b}}$ is distributed as $(\beta)/\mathfrak{b}$, which finishes the proof. \square

Lemma 6.23 (Lifting property of distributions). *Suppose that a distribution $\mathcal{D} : \text{Div}_{K^m}^0 \rightarrow \mathbb{R}$ satisfies $\|[\mathcal{D}] - \mathcal{U}(\text{Pic}_{K^m}^0)\|_1 < \varepsilon$ (see Definition 6.8). Then there exists a ‘lifted’ distribution $\mathcal{D}_U : \text{Div}_K^0 \rightarrow \mathbb{R}^+$ such that $[\mathcal{D}_U] = \mathcal{U}(\text{Pic}_{K^m}^0)$ and $\|\mathcal{D} - \mathcal{D}_U\|_1 < \varepsilon$.*

Proof. Put

$$\mathcal{D}_U(\mathbf{a}) = \begin{cases} \frac{1}{\text{Vol}(\text{Pic}_{K^m}^0)} \cdot \frac{\mathcal{D}(\mathbf{a})}{[\mathcal{D}]([\mathbf{a}])} & \text{if } [\mathcal{D}]([\mathbf{a}]) \neq 0 \\ u & \text{otherwise} \end{cases},$$

for some $u : \text{Div}_{K^m}^0 \rightarrow \mathbb{R}^+$ that satisfies $[u] = \frac{1}{\text{Vol}(\text{Pic}_{K^m}^0)}$. Then, one can check that $[\mathcal{D}_U] = \frac{1}{\text{Vol}(\text{Pic}_{K^m}^0)}$ is uniform on $\text{Pic}_{K^m}^0$. Furthermore, writing F for a fundamental domain in $\text{Div}_{K^m}^0$ for $\text{Pic}_{K^m}^0$, we have

$$\begin{aligned} \|\mathcal{D} - \mathcal{D}_U\|_1 &= \int_{\mathbf{a} \in F} \int_{\alpha \in K^*/\mu_K} |\mathcal{D}(\mathbf{a} + \langle \alpha \rangle) - \mathcal{D}_U(\mathbf{a} + \langle \alpha \rangle)| d\alpha d\mathbf{a} \\ &= \int_{[\mathbf{a}] \in \text{Pic}_{K^m}^0} \left| [\mathcal{D}]([\mathbf{a}]) - \frac{1}{\text{Vol}(\text{Pic}_{K^m}^0)} \right| d([\mathbf{a}]) \\ &= \|[\mathcal{D}] - \mathcal{U}(\text{Pic}_{K^m}^0)\|_1 \leq \varepsilon. \end{aligned}$$

The first equation holds by definition, the second equation by the fact that the sign of $(\mathcal{D}(\mathbf{a} + \langle \alpha \rangle) - \mathcal{D}_U(\mathbf{a} + \langle \alpha \rangle))$ depends per construction solely on the coset $[\mathbf{a}]$. \square

7. The Power Residue Symbol is in ZPP

7.1. Summary

In this chapter we show that, assuming the Riemann hypothesis for Hecke L-functions on the cyclotomic fields $\mathbb{Q}(\zeta_m)$, the problem of computing the m -th power residue symbol in a field containing the m -th root of unity lies in the complexity class ZPP. In other words, there exists an algorithm that computes power residue symbols within probabilistic polynomial time in the input size. Though this algorithm never outputs an incorrect output, it might simply give *no* output with a certain constant probability. The probability here is over, say, random coin flips, which allows the algorithm to repeat until having a negligible error probability. Such algorithms are also known as Las Vegas algorithms.

The proof of the polynomial running time consists of essentially two parts, which are treated separately in Section 7.4 and Section 7.5. The former part consists of an efficient reduction from general power residue symbols to power residue symbols in cyclotomic fields; this reduction is due to Lenstra [Len95] and Squirrel [Squ97]. The latter part is a new result and consists of a proof that power residue symbols in cyclotomic fields can be computed efficiently, assuming the Extended Riemann Hypothesis for Hecke L-functions on cyclotomic fields. The key ingredient for this algorithm to be provable is the sampling algorithm of the previous Chapter 6. By combining these two parts, one obtains a conditional proof that power residue symbols can be computed efficiently in any number field.

7.2. Introduction

The power residue symbol often plays a significant role in algorithms in which *residuosity* is involved, which is about distinguishing m -th powers from non- m -th powers modulo an ideal in a number field. In such case, the m -th power residue symbol serves as a first check, as it should be equal to one in the case of an m -th power.

Examples of cryptographic schemes involving residuosity and that need a fast computation of power residue symbols include [SW95; GM84; Sch98; Wil85], which mostly consider m being prime and below 12. It should be noted that these cryptographic schemes (and actually, most residuosity-based schemes) are not quantum secure, due to their susceptibility to Shor's efficient quantum algorithm for factoring [Sho94]. In fact, if one is allowed to use a quantum computer, a very simple algorithm for the power residue exists, by just factoring the bottom input ideal of $(\frac{\alpha}{b})$. So, to be clear, in this chapter we will solely consider classical computing power.

Efficient algorithms for the m -th power residue symbol for specific small cases of $m \leq 11$ are studied extensively [CS10; DF05; Wei02; Wil85; SW95; Lem00]. A first attempt to design an efficient algorithm for *general* m -th power residue symbols (i.e., for all m) was done by Squirrel in his undergraduate thesis [Squ97]. In that work, Squirrel derives an efficient reduction from power residue symbols in general number fields to those in cyclotomic fields based on an idea of Lenstra [Len95]. Squirrel also proposes an algorithm for computing power residue symbols in cyclotomic fields, but it relies on heavy precomputations and is therefore not polynomial for varying m [Squ97, Ch. 5, §3]. On top of that, the algorithm also seems unfeasible in terms of practical running time.

Later, an algorithm for m -th power residue symbols that seems practically feasible and runs heuristically in polynomial time (for varying m) was given by the author of this PhD thesis [Boe16; BP17]. In this chapter we prove that a variant of this heuristic algorithm lies in the complexity class ZPP, assuming the Extended Riemann Hypothesis for Hecke L-functions on cyclotomic fields.

It should be noted that the aforementioned algorithms tailored to specific small $m \leq 11$ are by far more efficient than this more general algorithm, are mostly deterministic and also do not require any variant of the Riemann hypothesis.

Difference between the power residue symbol algorithm of this chapter and the heuristic algorithm in [Boe16; BP17]

The key difference between the power residue symbol algorithm of this chapter and that of [Boe16; BP17] is their *purpose*. The algorithm described in this chapter is namely specifically constructed in such a way that the proof of its polynomial time complexity is as simple as possible. The heuristic algorithm in [Boe16; BP17], however, is much more directed toward implementation and a fast practical running time (for an implementation, see [Boe17]). This distinction in purpose lead to the following key differences between the two algorithms.

The provable algorithm does not use the Hilbert reciprocity law. As opposed to the heuristic algorithm, the provable algorithm of this chapter does *not* use *Hilbert reciprocity*. In other words, the following reciprocity law involving Hilbert symbols does not play any role in the provable algorithm of this chapter.

$$\left(\frac{\alpha}{\beta}\right)_m \left(\frac{\beta}{\alpha}\right)_m^{-1} = \prod_{\mathfrak{p} \mid m\infty} (\alpha, \beta)_{\mathfrak{p}},$$

Here, $(\alpha, \beta)_{\mathfrak{p}}$ are the m -th Hilbert symbols in the completion $\mathbb{Q}(\zeta_m)_{\mathfrak{p}}$ (e.g., [Neu85, Ch. III, §5 and Ch. IV, §9]). Avoiding the Hilbert reciprocity law has as an advantage that there is no need to compute Hilbert symbols in the provable algorithm. In the heuristic algorithm, the computation of such Hilbert symbols relied on a efficient and provable algorithm of Bouw [Bou21].

The provable algorithm uses the Artin reciprocity law. Instead, the provable algorithm of this chapter uses a different reciprocity law, namely the *Artin reciprocity* law (see Lemma 7.3), which states that for elements $\kappa \in K^{\mathfrak{m},1}$ in a specific *ray*, the power residue symbol $(\frac{\alpha}{\kappa}) = 1$ for all $\alpha \in K^*$. This turned out to be easier to use in a proof and has as an additional advantage that no computation of Hilbert symbols is needed. In fact, one can even use this provable algorithm to compute Hilbert symbols instead (see Section 7.6.1).

The provable algorithm does not use LLL-reduction. The heuristic algorithm of [Boe16; BP17] uses LLL-reduction to minimize sizes of the input while this is omitted in the provable algorithm for the sake of brevity and provability.

7.3. Preliminaries

In this chapter, K is a degree $n = [K : \mathbb{Q}]$ number field containing the m -th cyclotomic number field, i.e., $K \supseteq \mathbb{Q}(\zeta_m)$, where ζ_m is a primitive m -th root of unity. The main subject of this chapter is the *power residue symbol*, a map that partially captures m -th residuosity.

This power residue symbol takes as an input an ideal \mathfrak{b} in an order R of K and an element $\alpha \in R$, and outputs an m -th root of unity ζ_m^k . The symbol and its definition resembles that of the Jacobi symbol, for example in the sense that it can be defined in terms of prime ideals first, and can subsequently be multiplicatively extended to general ideals.

Definition 7.1 (Power residue symbol). *Let $\mathfrak{p} \nmid m$ be a prime ideal in an order R of $K \ni \zeta_m$ and let $\alpha \in R$ be an element coprime with m and \mathfrak{p} . We define $(\frac{\alpha}{\mathfrak{p}}) \in \langle \zeta_m \rangle = \{ \zeta_m^k \mid k \in \mathbb{N} \}$ to be the m -th root of unity that satisfies*

$$\left(\frac{\alpha}{\mathfrak{p}} \right) \equiv \alpha^{\frac{N(\mathfrak{p})-1}{m}} \pmod{\mathfrak{p}}.$$

For general ideals \mathfrak{b} in R coprime with m we then use the prime ideal factorization $\mathfrak{b} = \prod_j \mathfrak{p}_j^{e_j}$ to define the power residue symbol $(\frac{\alpha}{\mathfrak{b}})$.

$$\left(\frac{\alpha}{\mathfrak{b}}\right) := \prod_j \left(\frac{\alpha}{\mathfrak{p}_j}\right)^{e_j}.$$

By the very definition of the power residue symbol ‘above’ prime ideals, they can be computed efficiently and deterministically.

Lemma 7.2. *Let $\mathfrak{p} \subseteq \mathbb{Z}[\zeta_m]$ be a prime ideal not dividing m . Then the power residue symbol $(\frac{\alpha}{\mathfrak{p}})$ can be computed within $\text{poly}(m, \log \mathcal{N}(\mathfrak{p}), \log |\mathcal{N}(\alpha)|)$ time.*

Proof. By the power residue symbol formula for prime ideals we have $\left(\frac{\alpha}{\mathfrak{p}}\right) \equiv \alpha^{(\mathcal{N}(\mathfrak{p})-1)/m}$ modulo \mathfrak{p} . We compute the (modular) Hermite normal form [SL96; HM91] [Coh93, §2.4.2] of the ideal \mathfrak{p} , which allows to have a unique representative for each element in $\mathbb{Z}[\zeta_m]/\mathfrak{p}$. By modular exponentiation, can compute $\alpha^{(\mathcal{N}(\mathfrak{p})-1)/m}$ modulo \mathfrak{p} within time $\text{poly}(m, \log \mathcal{N}(\mathfrak{p}), \log |\mathcal{N}(\alpha)|)$. \square

The following lemma shows that the power residue symbol is trivial for certain values of the lower input. Specifically, considering a fixed upper input for the power residue symbol, the map $(\frac{\alpha}{\cdot}) : K \rightarrow \langle \zeta_m \rangle$ has a kernel that includes the *ray* $K^{\mathfrak{m}, 1}$ with $\mathfrak{m} = m^m \cdot \alpha$. This particular fact forms one of the very key ingredients of the efficient power residue symbol algorithm.

Lemma 7.3. *For all $\alpha \in \mathbb{Z}[\zeta_m]$ coprime with m , and all $\kappa \in \mathbb{Q}(\zeta_m)^*$ with $\text{ord}_{\mathfrak{p}}(\kappa) \geq 0$ for all $\mathfrak{p} \mid \alpha m$, we have,*

$$\left(\frac{\alpha}{1 + \kappa \cdot m^m \cdot \alpha}\right) = 1$$

Proof. Denote $K = \mathbb{Q}(\zeta_m)$ and $L = \mathbb{Q}(\zeta_m, \sqrt[m]{\alpha})$ for $\alpha \in \mathbb{Q}(\zeta_m)$. The power residue symbol $(\frac{\alpha}{b}) \in \langle \zeta_m \rangle$ in $\mathbb{Q}(\zeta_m)$ has the following relation with the Artin symbol [Lem00, §4.1] [Koc97, Ch. 2, §2.1]

$$\left(\frac{\alpha}{b}\right) \cdot \sqrt[m]{\alpha} = \left(\frac{b}{L/K}\right) [\sqrt[m]{\alpha}].$$

Denote $f_{L/K}$ for the *conductor* of the extension L/K . For any modulus m satisfying $f_{L/K} \mid m$, the *kernel* of the Artin symbol $\left(\frac{\cdot}{L/K}\right) : \mathcal{I}_K^m \rightarrow G$ contains the *ray* $K^{m,1}$, the multiplicative subgroup of K^* generated by elements $\kappa \in \mathbb{Z}[\zeta_m]$ that are 1 modulo m . This is a consequence of the Artin reciprocity law [Chi08, Thm. 2.1].

It remains to show that $m = m^m \alpha$ satisfies $f_{L/K} \mid m$, i.e., that $f_{L/K} \mid m^m \alpha$. If we can prove that fact, the result follows, since $1 + k \cdot m^m \cdot \alpha \in K^{m,1}$ for κ satisfying $\text{ord}_p(\kappa) \geq 0$ for all $p \mid \alpha m$.

In the following, we prove that $f_{L/K} \mid (m^m \alpha)$. Since α is required to be coprime with m , and the degree of the extension satisfies $[L : K] \mid m$, any $p \mid (\alpha)$ is tamely ramified in the extension L/K , because $p \nmid m$. Therefore, we have, [CG05, Ch. 2, Prop. 1.6.3] [CS08, Eq. (3.10) and Eq. (3.11)]

$$\text{For all } p \mid (\alpha) : \text{ord}_p(f_{L/K}) = 1.$$

From the same results, or from the fact that $f_{L/K} \mid \Delta_{L/K} \mid m^m \alpha^{m-1}$ [CF10, Lm. 5, Ch. 3] [NS13, Ch. VII, Prop. 11.9] follows that $f_{L/K} \mid (m^m \alpha)$. \square

The following result, namely, multiplicativity in the bottom input of the power residue symbol, can be found in [Neu85, Ch. 4, Eq. (9.2)] or [Koc97, Thm. 2.13].

Lemma 7.4. *Let K be a number field containing $\mathbb{Q}(\zeta_m)$. Let $b, c \in \mathcal{I}_K$ be coprime with m . For all $\alpha \in K$ coprime with b, c and m , we have*

$$\left(\frac{\alpha}{bc}\right) = \left(\frac{\alpha}{b}\right) \left(\frac{\alpha}{c}\right).$$

The last important lemma of this preliminaries concerns the local density of the prime ideals coprime to \mathfrak{m} . It turns out that for large enough r and not too large modulus \mathfrak{m} , the density $\delta_{\mathcal{P}^{\mathfrak{m}}}[r^n]$ does not differ so much from the density $\delta_{\mathcal{P}}[r^n]$ of all prime ideals in a number field. This density is known to be close to $\frac{1}{\rho_K \cdot \log(r^n)}$, where ρ_K is the residue of the Dedekind zeta function $\zeta_K(s)$ at the pole at $s = 1$.

This density is important because it is tightly related to the success probability of the power residue symbol algorithm of this chapter. This is because the power residue symbol algorithm involves prime ideal sampling, as in Chapter 6.

Lemma 7.5. *Let $\mathcal{P}^{\mathfrak{m}} = \{\mathfrak{p} \in \mathcal{I}_K^{\mathfrak{m}} \mid \mathfrak{p} \text{ prime}\}$ and let $\omega(\mathfrak{m})$ denote the number of different prime ideal divisors of \mathfrak{m} . Then, for all $r^n \geq \max((12 \log |\Delta_K| + 8n + 28)^4, 3 \cdot 10^{11}, 16 \cdot \omega(\mathfrak{m})^2)$, we have*

$$\delta_{\mathcal{P}^{\mathfrak{m}}}[r^n] \geq \frac{1}{4n \cdot \rho_K \cdot \log r}.$$

Recall that $n = [K : \mathbb{Q}]$, the degree of the number field K .

Proof. By Lemma 2.13, considering $x \in [(r/e)^n, r^n]$ and Definition 6.6, we have

$$\delta_{\mathcal{P}^{\mathfrak{m}}}[r^n] = \min_{x \in [(r/e)^n, r^n]} \frac{\pi_K^{\mathfrak{m}}(x)}{\rho_K x} \geq \frac{x / \log x}{4\rho_K x} \geq \frac{1}{4n\rho_K \log(r/e)} \geq \frac{1}{4n\rho_K \log r}.$$

□

7.4. Reduction to Cyclotomic Fields

7.4.1. Introduction

In this section, we will show that the computation of the m -th power residue symbol in any order R (of a number field) containing $\mathbb{Z}[\zeta_m]$ reduces to the

computation of (polynomially) many power residue symbols in $\mathbb{Z}[\zeta_m]$, the ring of integers of the m -th cyclotomic field.

The strategy of this proof is described in a paper of Lenstra [Len95], in which the special case $m = 2$ is elaborately worked out. For general $m > 2$, a full description of this reduction is given by Squirrel in his undergraduate thesis [Squ97]. In this section we will follow closely the reasoning of Squirrel and Lenstra, omitting precise complexity claims; any of the steps in this reduction runs in time polynomial in the input size.

In the following section, we give an overview of the proof of this reduction, postponing the definitions and proofs to a later moment.

7.4.2. Proof Strategy

Introduction. In this proof summary, we will consider number fields K containing all m -th roots of unity, i.e. $K \supseteq \mathbb{Q}(\zeta_m)$. Instead of the maximal order \mathcal{O}_K , which might be very hard to compute, we will mainly consider general orders $R \subseteq \mathcal{O}_K$ of K .

The main purpose of this proof overview is to show on a high level that we can reduce the computation of the power residue symbol $(\frac{\alpha}{\mathfrak{b}})_{m,K}$ for an element $\alpha \in R$ and an ideal $\mathfrak{b} \subseteq R$ to the computation of power residue symbols in the cyclotomic field $\mathbb{Q}(\zeta_m)$.

Signature identity. The power residue symbol $(\frac{\alpha}{\mathfrak{b}})_{m,K}$ is equal to another special quantity, $(m_\alpha, R/\mathfrak{b})$, which we will call the *signature*. This signature captures certain behavior of the multiplication map $m_\alpha : x \mapsto \alpha \cdot x$ on the finite $\mathbb{Z}[\zeta_m]$ -module R/\mathfrak{b} . Because of this equality, we can shift our attention to computing the signature $(m_\alpha, R/\mathfrak{b})$.

Invariant factor decomposition of R/\mathfrak{b} . A very important observation is the fact that the signature $(m_\alpha, R/\mathfrak{b})$ only depends on the structure of R/\mathfrak{b} as a $\mathbb{Z}[\zeta_m]$ -module. Using an analogue of the invariant factor decomposition

for finite modules over Dedekind domains (see [Coh99, Thm. 1.2.30]), we can efficiently compute a decomposition of R/\mathfrak{b} of the following shape.

$$R/\mathfrak{b} = \gamma_1 \mathbb{Z}[\zeta_m]/\mathfrak{d}_1 \oplus \cdots \oplus \gamma_k \mathbb{Z}[\zeta_m]/\mathfrak{d}_k, \quad (7.107)$$

where $\gamma \in R$ and \mathfrak{d}_j are ideals of $\mathbb{Z}[\zeta_m]$ that satisfy $\mathfrak{d}_j = \prod_{i=1}^j \mathfrak{c}_i$ for ideals \mathfrak{c}_i of $\mathbb{Z}[\zeta_m]$ that are neither the zero or the unit ideal. In other words, $\mathfrak{d}_{j+1}/\mathfrak{d}_j = \mathfrak{c}_{j+1}$ for $j \in \{1, \dots, k-1\}$ and $\mathfrak{d}_1 = \mathfrak{c}_1$. This computation shows that we can shift our focus to modules of a form as described in Equation (7.107).

The signature is compatible with short exact sequences. Let M', M, M'' be $\mathbb{Z}[\zeta_m]$ -modules with respective ($\mathbb{Z}[\zeta_m]$ -module compatible) automorphisms ϕ', ϕ and ϕ'' , that fit into the following commuting diagram

$$\begin{array}{ccccccc} 0 & \longrightarrow & M' & \longrightarrow & M & \longrightarrow & M'' \longrightarrow 0 \\ & & \downarrow \phi' & & \downarrow \phi & & \downarrow \phi'' \\ 0 & \longrightarrow & M' & \longrightarrow & M & \longrightarrow & M'' \longrightarrow 0 \end{array}$$

Then we have $(\phi, M) = (\phi', M') \cdot (\phi'', M'')$, i.e., the signature of the ‘middle’ module can be computed with the signatures of the ‘outer’ modules.

The determinant formula. For $\mathbb{Z}[\zeta_m]$ -modules isomorphic to $(\mathbb{Z}[\zeta_m]/\mathfrak{c})^t$ for some $t \in \mathbb{Z}_{>0}$ and integral ideal \mathfrak{c} of $\mathbb{Z}[\zeta_m]$, we can compute the signature by means of the determinant formula. Any automorphism ϕ of $(\mathbb{Z}[\zeta_m]/\mathfrak{c})^t$ can be described by a non-degenerate matrix with entries in $\mathbb{Z}[\zeta_m]/\mathfrak{c}$, which makes $\det(\phi) \in \mathbb{Z}[\zeta_m]/\mathfrak{c}$ a well-defined quantity. The determinant formula then reads as follows.

$$(\phi, (\mathbb{Z}[\zeta_m]/\mathfrak{c})^t) = \left(\frac{\det(\phi)}{\mathfrak{c}} \right)_{m, \mathbb{Q}(\zeta_m)}. \quad (7.108)$$

Note that this reduces the computation of this specific signature to a power residue symbol in the cyclotomic field $\mathbb{Q}(\zeta_m)$.

Applying induction on the components of the module. Denoting $M = R/\mathfrak{b}$, we have the following exact sequence

$$0 \rightarrow M/(\mathfrak{c}_1 M) \rightarrow M \rightarrow \mathfrak{c}_1 M \rightarrow 0,$$

where \mathfrak{c}_1 is the first factor in the invariant factor decomposition (Equation (7.107)). Because of the compatibility of the signature with short exact sequences, it is enough to compute the signatures $(\phi_\alpha, M/(\mathfrak{c}_1 M))$ and $(\phi_\alpha, \mathfrak{c}_1 M)$.

The first module, $M/(\mathfrak{c}_1 M)$, can be shown to be isomorphic to $(\mathbb{Z}[\zeta_m]/\mathfrak{c}_1)^k$, and therefore the determinant formula applies (see Equation (7.108)).

The last module, $\mathfrak{c}_1 M$, can be shown to have less ‘components’ than M itself; $k - 1$ instead of k .

$$\mathfrak{c}_1 M = \bigoplus_{j=1}^{k-1} \gamma_j \mathbb{Z}[\zeta_m]/\tilde{\mathfrak{d}}_j,$$

where $\tilde{\mathfrak{d}}_j = \mathfrak{d}_j/\mathfrak{c}_1$, and where \mathfrak{d}_j are obtained from the invariant factor decomposition of $M = R/\mathfrak{b}$.

Conclusion. By induction, we can therefore conclude that the computation of $(\phi_\alpha, R/\mathfrak{b})$ reduces to k power residue symbols $\left(\frac{d_j}{\mathfrak{c}_j}\right)_{m, \mathbb{Q}(\zeta_m)}$ for $j \in \{1, \dots, k\}$ in the cyclotomic field $\mathbb{Q}(\zeta_m)$. Here, \mathfrak{c}_j are the invariant factors of the module R/\mathfrak{b} as a $\mathbb{Z}[\zeta_m]$ -module and $d_j \in \mathbb{Z}[\zeta_m]/\mathfrak{c}_j$ are determinants of associated automorphisms.

7.4.3. Signature Identity

Definition 7.6 (Admissible modules). *We call a $\mathbb{Z}[\zeta_m]$ -module M admissible if $|M|$ is finite and $\gcd(|M|, m) = 1$.*

Letting the group $\langle \zeta_m \rangle = \{\zeta_m^j \mid j \in \mathbb{Z}/m\mathbb{Z}\}$ act on an admissible module M , we can directly deduce that this action must be free on $M \setminus 0$. Namely, suppose

that there exists an $x \in M$ with $\zeta_m^j x = x$. Then we have $(\zeta_m^j - 1)x = 0$, which implies $mx = 0$ (as $(\zeta_m^j - 1) \mid m$). Since $|M|x = 0$, $mx = 0$ and $\gcd(|M|, m) = 1$, we have $1 \cdot x = 0$.

This means that $M \setminus 0$ can be written as a disjoint union of orbits $\langle \zeta_m \rangle \cdot x$ (for some $x \in M$), where the orbits have precisely m elements. This directly implies $|M| = tm + 1$, where t is the number of orbits in $M \setminus 0$. Summarizing, any admissible module M satisfies $|M| \equiv 1$ modulo m .

Let M be an admissible $\mathbb{Z}[\zeta_m]$ -module and let $\phi : M \rightarrow M$ be a bijective function satisfying $\phi(\zeta_m \cdot x) = \zeta_m \cdot \phi(x)$ for all $x \in M$. Then ϕ acts faithfully on the $\langle \zeta_m \rangle$ -orbits of M , as $\phi(\langle \zeta_m \rangle \cdot x) = \langle \zeta_m \rangle \cdot \phi(x)$. In other words, ϕ induces a permutation on the quotient set $M/\langle \zeta_m \rangle$, fixing $0 \in M$.

Example 7.7. Put $K = \mathbb{Q}(\zeta_6, \sqrt[3]{2})$, a degree 6 extension of \mathbb{Q} . The subring $R = \mathbb{Z}[\zeta_6, \sqrt[3]{2}]$ is an order in K which has the following R -ideal $\mathfrak{p}_5 = (5, 3 - \sqrt[3]{2})$. Then the $\mathbb{Z}[\zeta_6]$ -module R/\mathfrak{p}_5 has 25 elements; one of them is zero, and the others fall into four $\langle \zeta_6 \rangle$ -orbits of length six, see Figure 7.1.

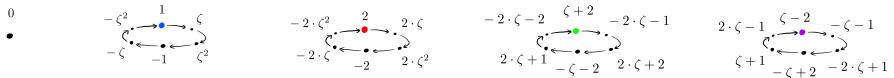


Figure 7.1.: The multiplicative action of $\langle \zeta \rangle$ on the 25 elements of R/\mathfrak{p}_5 as in Example 7.7, where $\zeta = \zeta_6$, a 6-th primitive root of unity. It consists of one zero-orbit of length one, and four orbits of length 6.

Let $S \subseteq M$ be a representative set for $M/\langle \zeta_m \rangle$, i.e., $M = \bigcup_{s \in S} \langle \zeta_m \rangle s$ (where the union is disjoint). Then, the action of ϕ on $M/\langle \zeta_m \rangle$ induces a bijection $s \mapsto s^\phi$ on S . Here $s^\phi \in S$ is the unique representative in S satisfying $\phi(\langle \zeta_m \rangle s) = \langle \zeta_m \rangle s^\phi$. Note that this means that $\phi(s) \in \langle \zeta_m \rangle s^\phi$, making the fraction $\frac{\phi(s)}{s^\phi} \in \langle \zeta_m \rangle$ well-defined for all $s \in S \setminus 0$. We then arrive at the following definition.

Definition 7.8 (Signature). *Let M be an admissible $\mathbb{Z}[\zeta_m]$ -module, let $\phi : M \rightarrow M$ be $\mathbb{Z}[\zeta_m]$ -module homomorphism and let $S \subseteq M$ be a representative*

7. The Power Residue Symbol is in ZPP

set for $M/\langle \zeta_m \rangle$. Then we define the signature $(\phi, M) \in \langle \zeta_m \rangle$ as follows.

$$(\phi, M) = \prod_{s \in S \setminus 0} \frac{\phi(s)}{s^\phi} \quad (7.109)$$

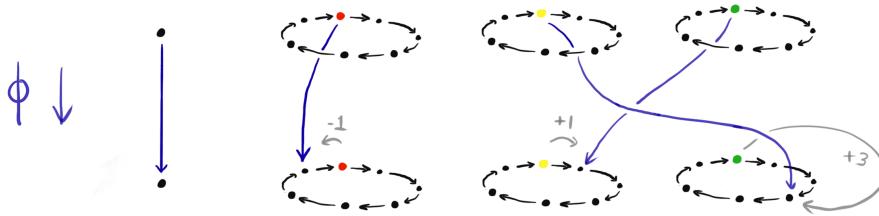


Figure 7.2.: The signature of a map ϕ forgets about the permutation of the $\langle \zeta_m \rangle$ -orbits. Instead, it captures the ‘compound deviation’ of the images of representatives from the representative of the orbits that image lives in. For example, the ϕ -image of the green dot deviates $+3$ from the yellow representative in its orbit.

Remark 7.9. The definition above can be generalized to any bijective map $M \rightarrow M$ that commutes with ζ_m [Squ97], but for our purposes it is enough to consider $\mathbb{Z}[\zeta_m]$ -module homomorphisms.

The very nature of the definition shows that (ϕ, M) does not depend on the choice of the representative set S . Namely, changing a single $s \in S$ into $s' = \zeta_m^j \cdot s$ causes a ζ_m^j to appear once in the numerator of a factor in Equation (7.109) and once in the denominator of a factor in Equation (7.109); therefore it does not change the overall value.

Lemma 7.10. Let R be an order in a number field K with $\mathbb{Z}[\zeta_m] \subseteq R$. Let \mathfrak{p} be a prime ideal in R , coprime with m . Let $\alpha \in R$ such that $\bar{\alpha} = \alpha \bmod \mathfrak{p} \in (R/\mathfrak{p})^*$ and denote $\phi_\alpha : R/\mathfrak{p} \rightarrow R/\mathfrak{p}, x \mapsto \bar{\alpha} \cdot x$. Then

$$\left(\frac{\alpha}{\mathfrak{p}} \right)_m = (\phi_\alpha, R/\mathfrak{p})$$

Proof. Taking a representative set S for $M = R/\mathfrak{p}$ (modulo $\langle \zeta_m \bmod \mathfrak{p} \rangle$) we write out the definition of $(\phi_\alpha, R/\mathfrak{p})$ (see Definition 7.8). In the following chain of equalities we make use of the fact that $M = R/\mathfrak{p}$ (next to a $\mathbb{Z}[\zeta_m]$ -module) is also a field, so that division and multiplication of elements there make sense.

$$(\phi_\alpha, R/\mathfrak{p}) = \prod_{s \in S \setminus 0} \frac{\phi_\alpha(s)}{s^{\phi_\alpha}} = \prod_{s \in S \setminus 0} \frac{\bar{\alpha} \cdot s}{s^{\phi_\alpha}} = \bar{\alpha}^{|S \setminus 0|} \frac{\prod_{s \in S \setminus 0} s}{\prod_{s \in S \setminus 0} s^{\phi_\alpha}} = \bar{\alpha}^{|S \setminus 0|}.$$

The last inequality follows from the fact that $s \mapsto s^{\phi_\alpha}$ is a bijection on $S \setminus 0$. As $|S \setminus 0| = \frac{|M|-1}{m} = \frac{\mathcal{N}(\mathfrak{p})-1}{m}$, we conclude that¹ $(\phi_\alpha, R/\mathfrak{p}) = \alpha^{(\mathcal{N}(\mathfrak{p})-1)/m} \bmod \mathfrak{p}$. This coincides with the definition of the power residue symbol $\left(\frac{\alpha}{\mathfrak{p}}\right)_m$. \square

Example 7.11. Put, again, $K = \mathbb{Q}(\zeta_6, \sqrt[3]{2})$ with order $R = \mathbb{Z}[\zeta_6, \sqrt[3]{2}]$ and the R -ideal $\mathfrak{p}_5 = (5, 3 - \sqrt[3]{2})$, as in Example 7.7. Putting $\alpha = \zeta_6 + 1$, we want to verify that $\left(\frac{\zeta_6+1}{\mathfrak{p}_5}\right)_6 = (\phi_{\zeta_6+1}, R/\mathfrak{p}_5)$, as in Lemma 7.10. The computation of $\left(\frac{\zeta_6+1}{\mathfrak{p}_5}\right)_6$ happens by observing that $\mathcal{N}(\mathfrak{p}_5) = 25$ and computing (using Lemma 7.2)

$$\begin{aligned} (\zeta_6 + 1)^{\frac{\mathcal{N}(\mathfrak{p}_5)-1}{6}} &= (\zeta_6 + 1)^4 = \zeta_6^4 + 4 \cdot \zeta_6^3 + 6 \cdot \zeta_6^2 + 4 \cdot \zeta_6 + 1 \\ &\equiv 9 \cdot \zeta_6 + 9 \equiv -(\zeta_6 + 1) = \zeta_6^5 \bmod \mathfrak{p}_5. \end{aligned}$$

Therefore, $\left(\frac{\alpha}{\mathfrak{p}_5}\right)_6 = \zeta_6^5$. The computation of the signature gives the same result, as can be seen in Figure 7.3. For the computation of the images in that figure; $\phi_{1+\zeta}(\zeta + 2) = (1 + \zeta)(2 + \zeta) = 4 \cdot \zeta + 1 \equiv -\zeta + 1 = -\zeta^2 \pmod{\mathfrak{p}_5}$ and $\phi_{\zeta+1}(\zeta - 2) = (1 + \zeta)(\zeta - 2) = -3 \equiv 2 \pmod{\mathfrak{p}_5}$.

For later purposes, we will need the following lemma, which shows that the signature map $(\cdot, M) : \text{Aut}_{\mathbb{Z}[\zeta_m]}(M) \rightarrow \langle \zeta_m \rangle$ is a group homomorphism.

Lemma 7.12. For two automorphisms ϕ, ψ of an admissible module M , we have

$$(\phi \circ \psi, M) = (\phi, M) \cdot (\psi, M)$$

¹Note that this element $\bar{\alpha}^{(\mathcal{N}(\mathfrak{p})-1)/m}$ coincides with the action of multiplication $x \mapsto \zeta_m^j x$ on R/\mathfrak{p} for some $j \in \mathbb{Z}/m\mathbb{Z}$.

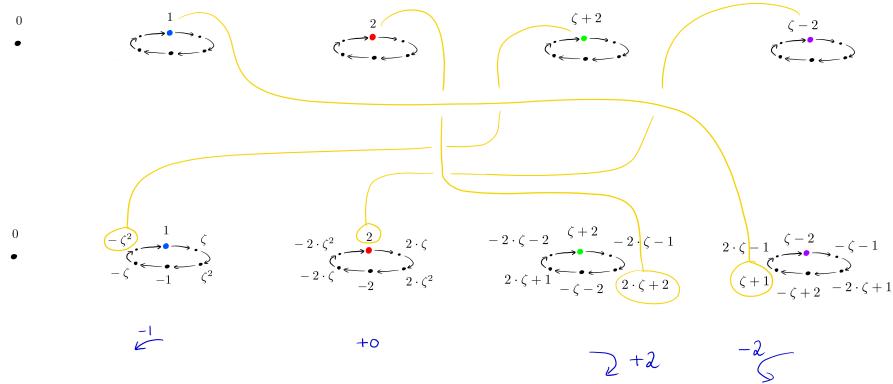


Figure 7.3.: The computation of the signature $(\phi_{1+\zeta}, R/\mathfrak{p}_5)$ of the map $\phi_{1+\zeta}(x)$, given by the rule $\phi_{1+\zeta}(x) = (1 + \zeta) \cdot x$, as in Example 7.11. By taking the sum of the images' displacements from the chosen representatives (the colored points), we obtain $-1 + 0 + 2 - 2 = -1$. Therefore, we conclude that $(\phi_{1+\zeta}, R/\mathfrak{p}_5) = \zeta_6^{-1} = \zeta_6^5$.

Proof. Choose a representative system S of $M/\langle \zeta_m \rangle$. Then

$$\begin{aligned}
 (\phi\psi, M) &= \prod_{s \in S} \frac{\phi(\psi(s))}{s^{\phi\psi}} = \prod_{s \in S} \frac{\phi(\psi(s))}{(s^\psi)^\phi} = \prod_{s \in S} \frac{\phi(\psi(s))}{\phi(s^\psi)} \frac{\phi(s^\psi)}{(s^\psi)^\phi} \\
 &= \phi \left(\prod_{s \in S} \frac{\psi(s)}{s^\psi} \right) \prod_{s \in S} \frac{\phi(s)}{s^\phi} = \phi((\psi, M)) \cdot (\phi, M) = (\psi, M) \cdot (\phi, M).
 \end{aligned}$$

□

7.4.4. Invariant Factor Decomposition of R/\mathfrak{b}

Computing the invariant factor decomposition of R/\mathfrak{b} as a module over $\mathbb{Z}[\zeta_m]$ happens by means of the Smith normal form in Dedekind domains (see [Coh99, §1.7]).

This particular Smith normal form algorithm as described in Cohen's book [Coh99, §1.7], needs modules to be represented in terms of *pseudobases*.

Usually, (in a computer algebra system) an R -ideal \mathfrak{b} is represented by means of a basis over \mathbb{Z} instead. We shortly describe here how to obtain such a pseudobasis from a \mathbb{Z} -basis. Let $\mathfrak{b} = \sum_{j=1}^t \mathbb{Z}\beta_j$. Then it is clear that the same set $(\beta_j)_{j \in \{1, \dots, t\}}$ is also a generating set over $\mathbb{Z}[\zeta_m]$, that is: $\mathfrak{b} = \sum_{j=1}^t \mathbb{Z}[\zeta_m]\beta_j$. By using the Hermite normal form over Dedekind domains [Coh99, §1.4] that removes linear dependencies, we arrive at a pseudobasis of \mathfrak{b} over $\mathbb{Z}[\zeta_m]$. The exact same reasoning can be applied to obtain a pseudobasis the ring R as a module over $\mathbb{Z}[\zeta_m]$.

Remark 7.13. *In the undergraduate thesis of Squirrel [Squ97], this step is partially done by computing $\mathbb{Z}[\zeta_m]$ -annihilators of the module R/\mathfrak{b} [Squ97, Ch. 4, §3].*

By [Coh99, §1.7], using a modular Smith normal form, we can deduce that we can find pseudobases for R and \mathfrak{b} of the following shape. $R = \bigoplus_{j=1}^t \mathfrak{s}_j \omega_j$, and $\mathfrak{b} = \bigoplus_{j=1}^t \mathfrak{d}_j \mathfrak{s}_j \omega_j$ where \mathfrak{s}_j are ideals of $\mathbb{Z}[\zeta_m]$, \mathfrak{d}_j are integral ideals of $\mathbb{Z}[\zeta_m]$ satisfying $\mathfrak{d}_{j-1} \subsetneq \mathfrak{d}_j$ for $j \geq 2$ and $\omega_j \in R$. This means that $R/\mathfrak{b} \xrightarrow{\sim} \bigoplus_{j=1}^t \mathbb{Z}[\zeta_m]/\mathfrak{d}_j$.

7.4.5. The Signature is Compatible with Short Exact Sequences

Proposition 7.14. *Let M', M, M'' be admissible $\mathbb{Z}[\zeta_m]$ -modules and let ϕ', ϕ, ϕ'' be $\mathbb{Z}[\zeta_m]$ -module automorphisms of M', M, M'' such that the following diagram commutes.*

$$\begin{array}{ccccccc} 0 & \longrightarrow & M' & \xrightarrow{\iota} & M & \xrightarrow{\pi} & M'' \longrightarrow 0 \\ & & \downarrow \phi' & & \downarrow \phi & & \downarrow \phi'' \\ 0 & \longrightarrow & M' & \xrightarrow{\iota} & M & \xrightarrow{\pi} & M'' \longrightarrow 0 \end{array}$$

Then

$$(\phi', M')(\phi'', M'') = (\phi, M).$$

Proof. Let S' be a representative set of $M'/\langle\zeta_m\rangle$. Extend $\iota(S')$ with a (disjoint) set $S'' \subseteq M$ such that $S = \iota(S') \cup S''$ is a representative set of $M/\langle\zeta_m\rangle$. Then

$$(\phi, M) = \prod_{s \in S} \frac{\phi(s)}{s^\phi} = \prod_{s' \in S'} \frac{\phi(\iota(s'))}{\iota(s')^\phi} \cdot \prod_{s'' \in S''} \frac{\phi(s'')}{(s'')^\phi} = (\phi', M')(\phi'', M''), \quad (7.110)$$

where the last equality is proven in two parts.

(i) As $\phi\iota = \iota\phi'$, we have $\iota(s')^\phi = \iota((s')^{\phi'})$. Therefore,

$$\prod_{s' \in S'} \frac{\phi(\iota(s'))}{\iota(s')^\phi} = \iota \left(\prod_{s' \in S'} \frac{\phi'(s')}{(s')^{\phi'}} \right) = \iota((\phi', M')) = (\phi', M').$$

(ii) Since S'' is distinct from $\iota(S')$, none of the $s'' \in S''$ send to zero under π . Therefore, we can apply π to the rightmost factor in Equation (7.110).

$$\pi \left(\prod_{s'' \in S''} \frac{\phi(s'')}{(s'')^\phi} \right) = \prod_{s'' \in S''} \frac{\pi\phi(s'')}{\pi((s'')^\phi)} = \prod_{s'' \in S''} \frac{\phi(\pi(s''))}{\pi(s'')^{\phi''}} \quad (7.111)$$

As S'' covers all $\langle\zeta_m\rangle$ -orbits of M that do not send to zero under π , the map $S'' \rightarrow \pi(S'')$, $s'' \mapsto \pi(s'')$ is a $|M'|$ -to-one map, i.e., $|\pi(S'')| = |S''|/|M'|$. Also, by surjectivity, $\pi(S'')$ is a representative set for the set $(M'' \setminus 0)/\langle\zeta_m\rangle$. Therefore, Equation (7.111) equals

$$\left(\prod_{t \in \pi(S'')} \frac{\phi''(t)}{t^{\phi''}} \right)^{|M'|} = ((\phi'', M''))^{|M'|} = (\phi'', M''),$$

where the last equality holds because $|M'| \equiv 1$ modulo m and $(\phi'', M'') \in \langle\zeta_m\rangle$.

□

Lemma 7.15. *Let R be an order in a number field K with $\mathbb{Z}[\zeta_m] \subseteq R$. Let \mathfrak{b} be an ideal in R , coprime with m . Let $\alpha \in R$ such that $\bar{\alpha} = \alpha \bmod \mathfrak{b} \in (R/\mathfrak{b})^*$ and denote $\phi_\alpha : R/\mathfrak{b} \rightarrow R/\mathfrak{b}, x \mapsto \bar{\alpha} \cdot x$. Then*

$$\left(\frac{\alpha}{\mathfrak{b}} \right)_m = (\phi_\alpha, R/\mathfrak{b})$$

Proof. We proceed by induction on the number of different prime ideal factors of \mathfrak{b} . The base case consists of \mathfrak{b} having only one prime divisor, i.e., $\mathfrak{b} = \mathfrak{p}^k$ being a prime power. If $k = 1$, we can apply Lemma 7.10. If $k > 1$, we can construct the following exact sequence

$$0 \rightarrow R/\mathfrak{p} \rightarrow R/\mathfrak{p}^k \rightarrow R/\mathfrak{p}^{k-1} \rightarrow 0$$

where the injection map is defined (non-canonically) by multiplying by an element $\gamma \in \mathfrak{p}^{k-1} \setminus \mathfrak{p}^k$. Then, together with the multiplication-by- α map (which we conveniently write ϕ_α for all rings involved), above exact sequence satisfies the conditions of Proposition 7.14. Therefore, by induction,

$$(\phi_\alpha, R/\mathfrak{p}^k) = (\phi_\alpha, R/\mathfrak{p}^{k-1}) \cdot (\phi_\alpha, R/\mathfrak{p}) = \left(\frac{\alpha}{\mathfrak{p}^{k-1}} \right)_m \left(\frac{\alpha}{\mathfrak{p}} \right)_m = \left(\frac{\alpha}{\mathfrak{p}^k} \right)_m.$$

The induction step consists of \mathfrak{b} being not a prime power. In that case, we write $\mathfrak{b} = \mathfrak{p}^k \mathfrak{c}$ with \mathfrak{p} prime, $k \geq 1$ and $\mathfrak{p} \nmid \mathfrak{c}$, and construct the following exact sequence

$$0 \rightarrow R/\mathfrak{p}^k \rightarrow R/\mathfrak{b} \rightarrow R/\mathfrak{c} \rightarrow 0,$$

where the injection $R/\mathfrak{p}^k \rightarrow \mathfrak{c}/\mathfrak{b}$ is defined (non-canonically) by multiplying by an element $\gamma \in \mathfrak{c}$ that satisfies $\gamma \equiv 1 \pmod{\mathfrak{p}^k}$. Again denoting ϕ_α for multiplication by α in all of the rings involved, this exact sequence satisfies the conditions of Proposition 7.14. Therefore, by induction,

$$(\phi_\alpha, R/\mathfrak{b}) = (\phi_\alpha, R/\mathfrak{p}^k) \cdot (\phi_\alpha, R/\mathfrak{c}) = \left(\frac{\alpha}{\mathfrak{p}^k} \right)_m \cdot \left(\frac{\alpha}{\mathfrak{c}} \right)_m = \left(\frac{\alpha}{\mathfrak{b}} \right)_m.$$

□

7.4.6. The Determinant Formula

Let $M = (\mathbb{Z}[\zeta_m]/\mathfrak{c})^t$ for some ideal \mathfrak{c} of $\mathbb{Z}[\zeta_m]$ and some $t \in \mathbb{N}_{>0}$. Then any automorphism $\phi : M \rightarrow M$ can be described as a non-degenerate $t \times t$ matrix with coefficients in $\mathbb{Z}[\zeta_m]/\mathfrak{c}$, which we call M_ϕ .

Lemma 7.16. *We have*

$$(\phi, (\mathbb{Z}[\zeta_m]/\mathfrak{c})^t) = \left(\frac{\det(M_\phi)}{\mathfrak{c}} \right)_{m, \mathbb{Q}(\zeta_m)}$$

Proof. We prove the statement first for $\mathfrak{c} = \mathfrak{p}$ a prime ideal. In that case the matrix M_ϕ has coefficients in the field $\mathbb{Z}[\zeta_m]/\mathfrak{p}$. Matrices over fields can be decomposed into $M_\phi = ULU'$, where U, U' are upper triangular and L is lower triangular, by means of Gaussian elimination. We denote ϕ_U, ϕ_L, ϕ'_U for their associated maps on $(\mathbb{Z}[\zeta_m]/\mathfrak{p})^t$. We have the exact sequence

$$0 \rightarrow \mathbb{Z}[\zeta_m]/\mathfrak{p} \rightarrow (\mathbb{Z}[\zeta_m]/\mathfrak{p})^t \rightarrow (\mathbb{Z}[\zeta_m]/\mathfrak{p})^{t-1} \rightarrow 0$$

where the injection map is just $x \mapsto (x, 0, \dots, 0)$ and the projection map projects on the last $t-1$ coordinates. By the (upper/lower) triangular shape of the matrix U of ϕ_U and by induction, one can deduce that

$$(\phi_U, M) = \left(\frac{\det(U)}{\mathfrak{p}} \right),$$

and the same for U' and L . Therefore,

$$\begin{aligned} (\phi, M) &= (\phi_U \phi_L \phi_{U'}, M) = (\phi_U, M)(\phi_L, M)(\phi_{U'}, M) \\ &= \left(\frac{\det(U)}{\mathfrak{p}} \right) \left(\frac{\det(L)}{\mathfrak{p}} \right) \left(\frac{\det(U')}{\mathfrak{p}} \right) = \left(\frac{\det(ULU')}{\mathfrak{p}} \right) = \left(\frac{M_\phi}{\mathfrak{p}} \right). \end{aligned}$$

This proves the statement for \mathfrak{c} being a prime ideal. For the general case, write $\mathfrak{c} = \mathfrak{p}\mathfrak{a}$, and construct the exact sequence

$$0 \rightarrow (\mathbb{Z}[\zeta_m]/\mathfrak{a})^t \rightarrow (\mathbb{Z}[\zeta_m]/\mathfrak{c})^t \rightarrow (\mathbb{Z}[\zeta_m]/\mathfrak{p})^t \rightarrow 0 \quad (7.112)$$

where the injection map is defined by scalar multiplication by $\varpi \in \mathfrak{p} \setminus \mathfrak{p}^2$ and the projection map just takes the entries modulo \mathfrak{p} .

Let $\phi' : (\mathbb{Z}[\zeta_m]/\mathfrak{a})^t \rightarrow (\mathbb{Z}[\zeta_m]/\mathfrak{a})^t$ respectively $\phi'' : (\mathbb{Z}[\zeta_m]/\mathfrak{p})^t \rightarrow (\mathbb{Z}[\zeta_m]/\mathfrak{p})^t$ be the map defined by reducing the entries of the matrix $M_\phi \in (\mathbb{Z}[\zeta_m]/\mathfrak{c})^{t \times t}$ modulo \mathfrak{a} respectively \mathfrak{p} . Then Equation (7.112) satisfies the requirements of Proposition 7.14, therefore

$$(\phi, (\mathbb{Z}[\zeta_m]/\mathfrak{c})^t) = (\phi', (\mathbb{Z}[\zeta_m]/\mathfrak{a})^t) \cdot (\phi'', (\mathbb{Z}[\zeta_m]/\mathfrak{p})^t)$$

$$= \left(\frac{\det(M_\phi)}{\mathfrak{a}} \right) \left(\frac{\det(M_\phi)}{\mathfrak{p}} \right) = \left(\frac{\det(M_\phi)}{\mathfrak{c}} \right).$$

Here, we used the induction hypothesis, the fact that

$$\det(M_{\phi'}) = \det(M_\phi \bmod \mathfrak{a}) = \det(M_\phi) \bmod \mathfrak{a},$$

and the similar statement for \mathfrak{p} . \square

7.4.7. Applying Induction on the Components of the Module

Lemma 7.17. *Let $R \subseteq K$ be a number ring containing a primitive m -th root of unity ζ_m and let $\mathfrak{b} \subseteq R$ be an ideal coprime with m . Let*

$$R/\mathfrak{b} = \gamma_1 \mathbb{Z}[\zeta_m]/\mathfrak{d}_1 \oplus \cdots \oplus \gamma_k \mathbb{Z}[\zeta_m]/\mathfrak{d}_k, \quad (7.113)$$

be the invariant factor decomposition of R/\mathfrak{b} with $\mathfrak{d}_j = \prod_{\ell \leq j} \mathfrak{c}_\ell$. Then we have, for all $\alpha \in R$ coprime with both \mathfrak{b} and m ,

$$\left(\frac{\alpha}{\mathfrak{b}} \right)_{m,K} = \prod_{j=1}^k \left(\frac{d_j}{\mathfrak{c}_j} \right)_{m, \mathbb{Q}(\zeta_m)},$$

where $d_j \in \mathbb{Z}[\zeta_m]/\mathfrak{c}_j$ are specific determinants of $k-j+1 \times k-j+1$ matrices with coefficients in $\mathbb{Z}[\zeta_m]/\mathfrak{c}_j$.

Proof. Denoting $M = R/\mathfrak{b}$, we have the following exact sequence

$$0 \rightarrow M/(\mathfrak{c}_1 M) \rightarrow M \rightarrow \mathfrak{c}_1 M \rightarrow 0,$$

where \mathfrak{c}_1 is the first factor in the invariant factor decomposition (Equation (7.107)). Because of the compatibility of the signature with short exact sequences, it is enough to compute the signatures $(m_\alpha, M/(\mathfrak{c}_1 M))$ and $(m_\alpha, \mathfrak{c}_1 M)$.

The first module, $M/(\mathfrak{c}_1 M)$, can be shown to be isomorphic to $(\mathbb{Z}[\zeta_m]/\mathfrak{c}_1)^k$, and therefore the determinant formula applies (see Equation (7.108)).

The last module, $\mathfrak{c}_1 M$, can be shown to have less ‘components’ than M itself; $k - 1$ instead of k .

$$\mathfrak{c}_1 M = \bigoplus_{j=1}^{k-1} \gamma_j \mathbb{Z}[\zeta_m] / \tilde{\mathfrak{d}}_j,$$

where $\tilde{\mathfrak{d}}_j = \mathfrak{d}_j / \mathfrak{c}_1$, and where \mathfrak{d}_j are obtained from the invariant factor decomposition of $M = R/\mathfrak{b}$. \square

7.4.8. Conclusion

By induction, we can therefore conclude that the computation of

$$\left(\frac{\alpha}{\mathfrak{b}} \right)_{m,R} = (\phi_\alpha, R/\mathfrak{b})$$

reduces to the computation of k power residue symbols $\left(\frac{d_j}{\mathfrak{c}_j} \right)_{m, \mathbb{Q}(\zeta_m)}$ (for $j \in \{1, \dots, k\}$) in the cyclotomic field $\mathbb{Q}(\zeta_m)$. Here, \mathfrak{c}_j are the invariant factors of the module R/\mathfrak{b} as a $\mathbb{Z}[\zeta_m]$ -module as in Equation (7.113) and $d_j \in \mathbb{Z}[\zeta_m]/\mathfrak{c}_j$ are determinants of associated automorphisms. We thus proved the following statement.

Theorem 7.18 (Lenstra, Squirrel). *Let $R \subseteq K$ be a number ring of a number field containing the m -th root of unity ζ_m . Let $\mathfrak{b} \subseteq R$ be an ideal coprime with m and let $\alpha \in R$ be an element of coprime with \mathfrak{b} and m . Then the computation of the power residue symbol $(\frac{\alpha}{\mathfrak{b}})_{m,R}$ reduces to at most $\log(\mathcal{N}(\mathfrak{b}))$ computations of the power residue symbols $\left(\frac{d_j}{\mathfrak{c}_j} \right)$ in $\mathbb{Q}(\zeta_m)$, where the d_j and \mathfrak{c}_j are bounded in size by the size of \mathfrak{b} and α .*

7.5. Computing the Power Residue Symbol in Cyclotomic Fields

7.5.1. Main Idea

Before explaining an algorithm in full detail, it is often insightful to give a simplified version first. The simplified version of the algorithm that computes power residue symbols $(\frac{\alpha}{\mathfrak{b}})$ with an element $\alpha \in \mathbb{Z}[\zeta_m]$ and an integral ideal \mathfrak{b} of $\mathbb{Z}[\zeta_m]$ essentially proceeds by two steps. An essential part of the algorithm is the idea that prime ideals ‘occur quite often’ in cyclotomic fields. This is a consequence of the density of primes of norm N being around $\frac{1}{\rho_K \log N}$ and the fact that the residue ρ_K of the Dedekind zeta function of cyclotomic fields at $s = 1$ is polynomially bounded (see Appendix A.2).

Step 1: Reducing the symbol $(\frac{\alpha}{\mathfrak{b}})$ to a ‘principal’ symbol $(\frac{\alpha}{\beta})$.

This happens by repeatedly sampling random $\beta \in \mathfrak{b}$ until the ideal $(\beta)/\mathfrak{b}$ is equal to some prime ideal \mathfrak{p} of $\mathbb{Z}[\zeta_m]$. In that case, write $(\beta) = \mathfrak{p}\mathfrak{b}$ and use the multiplicative property of the power residue symbol to obtain $(\frac{\alpha}{\beta}) = (\frac{\alpha}{\mathfrak{b}}) \cdot (\frac{\alpha}{\mathfrak{p}})$. By the fact that there exists an efficiently computable formula (see Lemma 7.2) for power residue symbols with a prime ideal as the bottom input, the symbol $(\frac{\alpha}{\mathfrak{p}})$ is efficiently computable.

Therefore, provided that such a suitable $\beta \in \mathfrak{b}$ can be efficiently found, the above procedure reduces the computation of the symbol $(\frac{\alpha}{\mathfrak{b}})$ to the computation of a power residue symbol $(\frac{\alpha}{\beta})$ where the bottom input β is an *element* in $\mathbb{Z}[\zeta_m]$ instead of a generic ideal.

Step 2: Evaluating the symbol $(\frac{\alpha}{\beta})$ by shifting β .

This happens by sampling random $\kappa \in \mathbb{Z}[\zeta_m]$ until the shifted element $\beta + \kappa m^m \alpha = \varpi$ is a *prime element*. As the power residue symbol $(\frac{\alpha}{\beta})$ with

$\alpha, \beta \in \mathbb{Z}[\zeta_m]$ satisfies the ‘shifting property’, (see Lemma 7.3) we have

$$\left(\frac{\alpha}{\beta}\right) = \left(\frac{\alpha}{\beta + \kappa m^m \alpha}\right) = \left(\frac{\alpha}{\varpi}\right).$$

Because (ϖ) is a prime ideal, there exists an efficiently computable formula for the symbol $(\frac{\alpha}{\varpi})$ (see Lemma 7.2). Therefore, $(\frac{\alpha}{\beta})$ can also be computed efficiently, provided that one indeed can find a $\kappa \in \mathbb{Z}[\zeta_m]$ such that $\beta + \kappa m^m \alpha$ is a prime element in $\mathbb{Z}[\zeta_m]$.

Discussion

It is clear that the first step only works whenever sampling a random $\beta \in \mathfrak{b}$ results sufficiently often in an ideal $(\beta)/\mathfrak{b}$ that is prime. In other words, the probability that $(\beta)/\mathfrak{b}$ is prime should be high enough. Likewise, the second step only works whenever sampling random $\kappa \in \mathbb{Z}[\zeta_m]$ results sufficiently often in an element $\beta + \kappa m^m \alpha$ that is prime.

It turns out to be notoriously hard to estimate these probabilities whenever \mathfrak{b} and β are *fixed*. However, if both \mathfrak{b} and β are appropriately *random* instead, one can actually lower bound these probabilities by means of *Landau’s prime ideal theorem*. This theorem can be informally expressed by saying that there are many prime ideals among the ideals in $\mathbb{Z}[\zeta_m]$. In other words, if one takes a ‘random ideal’ in $\mathbb{Z}[\zeta_m]$, there is a reasonable probability that it is a prime ideal.

So, in order to be fully able to estimate the success probability of the algorithm, we will need to appropriately *randomize* the lower input of the power residue symbol. With this adequate randomization, which will be done by means of a random walk as in Chapter 4 (thus relying on the Extended Riemann Hypothesis), one obtains the provable, full algorithm.

Remark 7.19. *In an actual implementation, one should not use this chapter’s provable algorithm. Instead, one should use the heuristic variant of it described in [BP17; Boe16]. A specific blend between the provable and the heuristic variant that uses Artin reciprocity (see Lemma 7.3) might also be*

Algorithm 8: POWERRESIDUESYMBOL(α, \mathfrak{b}, m), the computation of the symbol $(\frac{\alpha}{\mathfrak{b}})$

Require:

- An integer $m > 1$ defining the cyclotomic field $\mathbb{Q}(\zeta_m)$ of degree n .
- An integral element $\alpha \in \mathbb{Z}[\zeta_m]$ coprime with m .
- An integral ideal $\mathfrak{b} \subseteq \mathbb{Z}[\zeta_m]$ coprime with α and m ,

Ensure: $(\frac{\alpha}{\mathfrak{b}}) \in \langle \zeta_m \rangle$, or failure.

- 1: Put $\mathfrak{m} = m^m \cdot (\alpha)$ as the modulus.
- 2: Apply the sampling Algorithm 7 with \mathfrak{b} , \mathfrak{m} , $\tau = 1$ and $1/\varepsilon = \max(2^n, n^{5+1}(n + \log |\mathcal{N}(\alpha)|))$ to sample an element $\beta \in \tilde{\mathfrak{b}} \cap (1 + \mathfrak{m})$, where $\tilde{\mathfrak{b}} = \mathfrak{b} \prod_j \mathfrak{p}_j$ comes from the sampling algorithm.
- 3: **return** $\left(\frac{\alpha}{\mathfrak{p}}\right)^{-1} \cdot \prod_j \left(\frac{\alpha}{\mathfrak{p}_j}\right)^{-1}$ if $\beta/\tilde{\mathfrak{b}} = \mathfrak{p}$ is prime, using the formula for the power residue symbol above prime ideals (Lemma 7.2).
- 4: **return** *failure* otherwise.

interesting to implement, because it avoids the need for the computation of Hilbert symbols. Such an implementation (that relies on Artin reciprocity and not Hilbert reciprocity) might therefore even be used to compute Hilbert symbol due to a ‘global-to-local’ principle (see also Section 7.6.1).

7.5.2. The Full Algorithm

Lemma 7.20 (ERH). *Assuming the Riemann Hypothesis for Hecke L-functions on cyclotomic fields, Algorithm 8 is correct and runs in time polynomial in $m, \log |\mathcal{N}(\alpha)|$ and $\log \mathcal{N}(\mathfrak{b})$. Furthermore, Algorithm 8 has success probability at least*

$$\Omega\left(\frac{1}{n^5 \cdot (n^2 \log n + \log |\mathcal{N}(\alpha)|)}\right)$$

Proof. We start with proving the correctness of Algorithm 8, i.e., that the algorithm computes the symbol $(\frac{\alpha}{\tilde{b}})$ if it does not fail. This is proven by the sequence of equalities in Equation (7.114), which uses the multiplicative property of the power residue symbol (see Lemma 7.4) and the fact that the power residue symbol is trivial on the ray $K^{\mathfrak{m},1}$ with $\mathfrak{m} = m^m(\alpha)$ (see Lemma 7.3). So, since $\beta \in K^{\mathfrak{m},1}$ (i.e., $(\frac{\alpha}{\beta}) = 1$) and $(\beta) = \mathfrak{p}\tilde{\mathfrak{b}} = \mathfrak{p}\mathfrak{b}\prod_j \mathfrak{p}_j$, one obtains

$$1 = \left(\frac{\alpha}{\beta}\right) = \left(\frac{\alpha}{\mathfrak{p}}\right) \cdot \left(\frac{\alpha}{\tilde{\mathfrak{b}}}\right) = \left(\frac{\alpha}{\mathfrak{p}}\right) \cdot \left(\frac{\alpha}{\mathfrak{b}}\right) \cdot \prod_j \left(\frac{\alpha}{\mathfrak{p}_j}\right). \quad (7.114)$$

The correctness of the algorithm follows by rearranging terms to get an expression for $(\frac{\alpha}{\tilde{b}})$.

For the success probability, we need to estimate the probability that $(\beta)/\tilde{\mathfrak{b}}$ is a prime ideal in step 3. By the correspondence theorem between sampling probability and ideal density (see Theorem 6.21) we know that the probability of $(\beta)/\tilde{\mathfrak{b}}$ being prime equals at least $\delta_{\mathcal{S}^{\mathfrak{m}}}[r^n] - \varepsilon$, where $\mathcal{S}^{\mathfrak{m}} = \{\mathfrak{p} \in \mathcal{I}_K^{\mathfrak{m}} \mid \mathfrak{p} \text{ prime}\}$. By Lemma 7.5, the fact that $r^n \geq \mathcal{N}(\mathfrak{m}) \geq 16 \cdot \omega(\mathfrak{m})^2$, Writing out the instantiation for r in Algorithm 7, using $|\Delta_K|^{3/(2n)} \leq n^{3/2}$ for cyclotomic fields K , we have

$$\begin{aligned} r &= 4 \cdot 2^n \cdot n^{3/2} \cdot |\Delta_K|^{\frac{3}{2n}} \cdot \mathcal{N}(\mathfrak{m})^{1/n} \leq 4 \cdot 2^n \cdot n^3 \cdot \mathcal{N}(\mathfrak{m})^{1/n} \\ &\leq 2^{n+2} \cdot n^3 \cdot \mathcal{N}(m) \cdot \mathcal{N}(\alpha)^{1/n}, \end{aligned}$$

I.e., $\log(r^n) \leq n(n+2)\log(2) + 3n\log(n) + n^2\log n + \log |\mathcal{N}(\alpha)| = O(n^2 \log n + \log |\mathcal{N}(\alpha)|)$. Then, we have that the success probability is lower bounded (see Theorem 6.21) by

$$\delta_{\mathcal{S}^{\mathfrak{m}}}[r^n] - \varepsilon \geq \frac{1}{3 \cdot \rho_K \cdot n \cdot \log r} - \varepsilon \geq \frac{1}{\rho_K \cdot n \cdot (n^2 \log n + \log |\mathcal{N}(\alpha)|)} - \varepsilon$$

We show in Appendix A.2 that $\rho_K = O(n^4)$ (the hidden constant is $e^{15} \approx 3.3 \cdot 10^6$). By the instantiation $1/\varepsilon = \max(2^n, n^{5+1}(n^2 \log n + \log |\mathcal{N}(\alpha)|))$ we then have,

$$\begin{aligned} \delta_{\mathcal{S}^{\mathfrak{m}}}[r^n] - \varepsilon &= \Omega\left(\frac{1}{n^5 \cdot (n^2 \log n + \log |\mathcal{N}(\alpha)|)}\right) - \varepsilon \\ &= \Omega\left(\frac{1}{n^5 \cdot (n^2 \log n + \log |\mathcal{N}(\alpha)|)}\right) \end{aligned}$$

As Algorithm 7 is polynomial in its input size and in $\log(1/\varepsilon)$, it is enough to show that $\log \mathcal{N}(\mathfrak{p})$, $\log \mathcal{N}(\mathfrak{p}_j)$ and $\log(1/\varepsilon)$ are polynomially bounded in m , $\log |\mathcal{N}(\alpha)|$ and $\log \mathcal{N}(\mathfrak{b})$, in order to prove that Algorithm 8 runs in polynomial time.

Note that $\mathfrak{m} = m^m(\alpha)$, therefore $\log(\mathcal{N}(\mathfrak{m})) = \text{poly}(n, \log |\mathcal{N}(\alpha)|)$ is polynomially bounded. The logarithm of the inverse error $\log(1/\varepsilon)$ is easily shown to be polynomially bounded as well. Also $N, \log B$ and $\log r$ from Algorithm 7 with the instantiation of ε are polynomially bounded by $m, \log |\Delta_K| = O(m), \log(1/\varepsilon)$ and $\log \mathcal{N}(\mathfrak{d})$. So $\log \mathcal{N}(\mathfrak{p}_j) \leq \log B$ are polynomially bounded.

The largest prime, \mathfrak{p} , satisfies $\log \mathcal{N}(\mathfrak{p}) \leq \log(|\mathcal{N}(\beta)|/\mathcal{N}(\mathfrak{b})) \leq N \log B + n \log r$, by Algorithm 7. Therefore, all relevant quantities are polynomially bounded, thus the entire algorithm runs within polynomial time. \square

Theorem 7.21. *Let $K \supseteq \mathbb{Q}(\zeta_m)$ be a number field and let $R \subseteq K$ be an order in that number field. Assume the Extended Riemann Hypothesis for Hecke-L functions of the cyclotomic number field $\mathbb{Q}(\zeta_m)$.*

Then, there exists an algorithm that computes the power residue symbol $(\frac{\alpha}{\mathfrak{b}})$ for all elements $\alpha \in R$ and ideals $\mathfrak{b} \subseteq R$, within time polynomial in $\log |\Delta_K|$, $[K : \mathbb{Q}]$, $\text{size}(\alpha)$ and $\text{size}(\mathfrak{b})$.

Proof. Follows immediately from Lemma 7.20 and the reduction from Lenstra and Squirrel (Theorem 7.18). \square

7.6. Discussion

7.6.1. Computing Hilbert Symbols Using Power Residue Symbols

Because the algorithm in this chapter does not use the computation of Hilbert symbols (as opposed to the heuristic algorithm in [BP17; Boe16]), one can reverse the roles and use the computation of power residue symbols

to derive information about the associated Hilbert symbols in a number field K containing $\mathbb{Q}(\zeta_m)$ in the following way [Neu85, Ch. IV, §9].

$$\prod_{\mathfrak{p} \mid m\infty} (\alpha, \beta)_{\mathfrak{p}} = \left(\frac{\alpha}{\beta}\right)_m \left(\frac{\beta}{\alpha}\right)_m^{-1}.$$

To compute $(\alpha, \beta)_{\mathfrak{q}}$ for a fixed chosen $\mathfrak{q} \mid m$, one picks, using the Chinese remainder theorem, an element $\gamma \in \mathcal{O}_K$ that satisfies $\gamma \equiv 1$ modulo \mathfrak{p}^{d^2} for $\mathfrak{p} \mid m$ and $\mathfrak{p} \neq \mathfrak{q}$, and $\gamma \equiv \beta$ modulo \mathfrak{q}^{d^2} , where $d = [K : \mathbb{Q}]$ is the degree of the number field K . In that case, $(\alpha, \gamma)_{\mathfrak{p}} = 1$ for $\mathfrak{p} \neq \mathfrak{q}$ and $(\alpha, \gamma)_{\mathfrak{q}} = (\alpha, \beta)_{\mathfrak{q}}$, and therefore

$$(\alpha, \beta)_{\mathfrak{q}} = (\alpha, \gamma)_{\mathfrak{q}} = \prod_{\mathfrak{p} \mid m\infty} (\alpha, \gamma)_{\mathfrak{p}} = \left(\frac{\alpha}{\gamma}\right)_m \left(\frac{\gamma}{\alpha}\right)_m^{-1}.$$

In above reasoning, we use the following lemma.

Lemma 7.22. *Let $K_{\mathfrak{p}}$ be the completion of a number field $K \supseteq \mathbb{Q}(\zeta_m)$ of degree $d = [K : \mathbb{Q}]$ with respect to the finite prime $\mathfrak{p} \mid m$, and let $(\cdot, \cdot)_{\mathfrak{p}} : K_{\mathfrak{p}}^* \times K_{\mathfrak{p}}^* \rightarrow \langle \zeta_m \rangle$ denote the Hilbert symbol on this completion. Then*

$$(\alpha, 1 + \pi^{d^2})_{\mathfrak{p}} = 1 \text{ for all } \pi \in \mathfrak{p}.$$

Proof. As $(\alpha, \cdot)_{\mathfrak{p}} : K_{\mathfrak{p}} \rightarrow \langle \zeta_m \rangle$ equals the Artin symbol (or norm residue symbol) of the extension $K_{\mathfrak{p}}(\sqrt[m]{\alpha}) : K_{\mathfrak{p}}$ [Neu85, Ch. 3, Prop. 5.1], it suffices to show that $1 + \pi^{dm} \in \mathcal{N}_{K_{\mathfrak{p}}(\sqrt[m]{\alpha})/K_{\mathfrak{p}}}(K_{\mathfrak{p}}(\sqrt[m]{\alpha}))$ for all $\pi \in \mathfrak{p}$ [Neu85, Ch. 3, Prop. 5.2iii]. In other words, we need to show that the conductor $\mathfrak{f}_{K_{\mathfrak{p}}(\sqrt[m]{\alpha})/K_{\mathfrak{p}}}$ of this local Kummer extension divides \mathfrak{p}^{d^2} . By using local computations and Hensel's lemma [CS08, Eq. (3.11)], we know that

$$\text{ord}_{\mathfrak{p}}(\mathfrak{f}_{K_{\mathfrak{p}}(\sqrt[m]{\alpha})/K_{\mathfrak{p}}}) \leq d(1 + \log(d/m)) + 1 \leq d^2.$$

□

This leads to the following corollary.

Corollary 7.23. *Assuming the Extended Riemann Hypothesis for Hecke- L functions on cyclotomic fields, Hilbert symbols can be computed within time polynomial in the input size.*

This corollary is quite weak compared to the much stronger results of Bouw [Bou21]; his algorithm is unconditional (i.e., does not require any variant of the Riemann Hypothesis), deterministic, and his algorithm's studied complexity is way more explicit. Though no real-life comparison has been made yet, I suspect Bouw's algorithm to run significantly faster than the method described above.

7.6.2. Computing Artin Symbols in the Same Fashion

A very similar algorithm as Algorithm 8 could in principle be used to compute Artin symbols $\left(\frac{\cdot}{L/K}\right)$ for abelian extensions L/K . The main caveat is that the residue ρ_K of the Dedekind zeta function $\zeta_K(s)$ of K at $s = 1$ might be too large, i.e., not polynomially bounded. Such a large residue would make such an algorithm not feasible, as the success probability depends inversely on this residue ρ_K .

For the sake of completeness, we do spell out a proposal for an algorithm computing Artin symbols in Algorithm 9. We would like to stress that no guarantee on the running time is given, except maybe whenever the residue ρ_K is polynomially bounded. In that case, the proof resembles that of Lemma 7.20.

Remark 7.24. *To compute the Frobenius element $\left(\frac{\mathfrak{p}}{L/K}\right) \in G = \text{Gal}(L/K)$ for a prime \mathfrak{p} as in Line 3 of Algorithm 9, one goes through the following lines.*

- Compute $\mathfrak{P} \subseteq \mathcal{O}_L$, any prime ideal above $\mathfrak{p} \subseteq \mathcal{O}_K$.
- Compute a primitive element $\alpha \in L$, i.e., an $\alpha \in L$ such that $L = K(\alpha)$, by means of linear algebra.
- Compute $\alpha^q \bmod \mathfrak{P}$, where $q = |\mathcal{O}_K/\mathfrak{p}|$.
- Output a $g \in G = \text{Gal}(L/K)$ for which holds $\alpha^q \equiv g(\alpha) \bmod \mathfrak{P}$.

Algorithm 9: $\text{ArtinSymbol}(\mathfrak{b}, L, K)$, the computation of the Artin symbol $\left(\frac{\mathfrak{b}}{L/K}\right) \in \text{Gal}(L/K)$

Require:

- A number field extension L/K , where both L and K are defined by a defining polynomial over \mathbb{Q} .
- An integral ideal $\mathfrak{b} \subseteq \mathcal{O}_K$ coprime with Δ_L .
- For all $g \in G = \text{Gal}(L/K)$ and $\alpha \in L$, an efficient algorithm that computes $g(\alpha) \in L$.

Ensure: $\left(\frac{\alpha}{\mathfrak{b}}\right) \in \langle \zeta_m \rangle$, or failure.

- 1: Put $\mathfrak{m} = \Delta_{L/K}$ the relative discriminant of the extension L/K as the modulus.
- 2: Apply the sampling Algorithm 7 with \mathfrak{b} , \mathfrak{m} , $\tau = 1$ and $1/\varepsilon = \max(2^n, \rho_K \cdot n \cdot (n^2 \log n + n \log \mathcal{N}(\mathfrak{m})))$ to sample an element $\beta \in \tilde{\mathfrak{b}} \cap (1 + \mathfrak{m})$, where $\tilde{\mathfrak{b}} = \mathfrak{b} \prod_j \mathfrak{p}_j$ comes from the sampling algorithm.
- 3: **return** $\left(\frac{\mathfrak{p}}{L/K}\right)^{-1} \cdot \prod_j \left(\frac{\mathfrak{p}_j}{L/K}\right)^{-1}$ if $\beta/\tilde{\mathfrak{b}} = \mathfrak{p}$ is prime, using the formula for the Artin symbol for prime ideals ('Frobenius element', see [Neu85, Ch. IV, §8]).
- 4: **return** *failure* otherwise.

Remark 7.25. *The approach of Algorithm 9 is not expected to work for number fields with large Dedekind residue ρ_K . Though, we might enlarge the set of ‘good’ ideals \mathcal{S} by also including ‘near primes’, which are ideals that are a product of a large prime ideal and several smaller prime ideals; in other words, a large prime ideal times a smooth ideal.*

This might increase the local density of \mathcal{S} significantly in some cases, maybe even to the point that the Algorithm 9 succeeds within polynomial time even though ρ_K is not small.

An open question arising here is: What exactly does a large residue ρ_K mean? If it just implies more frequent small primes or more (higher) prime powers, it does not affect the Artin symbol algorithm. If it, on the other hand, implies a scarcity of easy-to-factor ideals, it does affect the Artin symbol algorithm. Are there means to distinguish these two cases?

Bibliography

[AC49] N. Ankeny and S. Chowla. “The class number of the cyclotomic field.” In: *Proceedings of the National Academy of Sciences of the United States of America* 35.9 (1949), pp. 529–532 (cit. on p. 290).

[Bac90] E. Bach. “Explicit bounds for primality testing and related problems.” In: *Mathematics of Computation* 55 (1990), pp. 355–380 (cit. on p. 162).

[Ban93] W. Banaszczyk. “New bounds in some transference theorems in the geometry of numbers.” In: *Mathematische Annalen* 296.4 (1993), pp. 625–636 (cit. on pp. 77, 86).

[BCP97] W. Bosma, J. Cannon, and C. Playoust. “The Magma algebra system. I. The user language.” In: *Journal of Symbolic Computation* 24.3-4 (1997), pp. 235–265 (cit. on p. 204).

[BDF20] K. de Boer, L. Ducas, and S. Fehr. “On the quantum complexity of the continuous hidden subgroup problem.” In: *EUROCRYPT*. Springer International Publishing, 2020, pp. 341–370 (cit. on p. 38).

[BF14] J.-F. Biasse and C. Fieker. “Subexponential class group and unit group computation in large degree number fields.” In: *LMS Journal of Computation and Mathematics* 17.A (2014), pp. 385–403 (cit. on pp. 134, 204–207).

[Bha+20] M. Bhargava, A. Shankar, T. Taniguchi, F. Thorne, J. Tsimerman, and Y. Zhao. “Bounds on 2-torsion in class groups of number fields and integral points on elliptic curves.” In: *Journal of the American Mathematical Society* 33.4 (Oct. 2020), pp. 1087–1099 (cit. on pp. 76, 282).

[Bia+17] J.-F. Biasse, T. Espitau, P.-A. Fouque, A. Gélin, and P. Kirchner. “Computing generator in cyclotomic integer rings.” In: *EUROCRYPT*. Springer. 2017, pp. 60–88 (cit. on p. 134).

[BK96] J. Buchmann and V. Kessler. “Computing a reduced lattice basis from a generating system.” In: *Unpublished Manuscript* (Aug. 1996) (cit. on pp. 83, 88, 95, 96, 124, 125).

[BKK17] L. Beilina, E. Karchevskii, and M. Karchevskii. *Numerical linear algebra: theory and applications*. Springer, Sept. 2017 (cit. on p. 126).

[Boe+20] K. de Boer, L. Ducas, A. Pellet-Mary, and B. Wesolowski. “Random self-reducibility of Ideal-SVP via Arakelov random walks.” In: *CRYPTO*. Springer International Publishing, 2020, pp. 243–273 (cit. on p. 39).

[Boe16] K. de Boer. “Computing the Power Residue Symbol.” Available at: http://koendeboer.com/publication/masterthesis/masterthesis_deBoer.pdf. MA thesis. Radboud Universiteit Nijmegen, The Netherlands, 2016 (cit. on pp. 236–238, 256, 259).

[Boe17] K. de Boer. *An implementation of the power residue symbol algorithm*. Available online at: <https://github.com/kodebro/powerresiduesymbol>. 2017 (cit. on p. 237).

[Bou21] J. Bouw. “On the computation of norm residue symbols.” PhD thesis. Universiteit Leiden, The Netherlands, 2021 (cit. on pp. 237, 261).

[BP17] K. de Boer and C. Pagano. “Calculating the power residue symbol and ibeta.” In: *ISSAC*. Vol. 68. 2017, pp. 923–934 (cit. on pp. 134, 204, 206, 236–238, 256, 259).

[BP89] J. Buchmann and M. Pohst. “Computing a lattice basis from a system of generating vectors.” In: *Proceedings of the European Conference on Computer Algebra. EUROCAL '87*. London, UK: Springer-Verlag, 1989, pp. 54–63 (cit. on pp. 95, 96, 123).

[BPR04] J. Buhler, C. Pomerance, and L. Robertson. “Heuristics for class numbers of prime-power real cyclotomic fields,” in: *High primes and misdemeanours: lectures in honour of the 60th birthday of Hugh Cowie Williams*. Fields Institute Communications. Amer. Math. Soc., 2004, pp. 149–157 (cit. on p. 55).

[Bre10] R. P. Brent. “Multiple-precision zero-finding methods and the complexity of elementary function evaluation.” In: *CoRR* abs /1004.3412 (2010) (cit. on p. 307).

[BS16] J.-F. Biasse and F. Song. “Efficient quantum algorithms for computing class groups and solving the principal ideal problem in arbitrary degree number fields.” In: *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM. 2016, pp. 893–902 (cit. on pp. 83, 86, 89, 166, 170).

[BS96] E. Bach and J. O. Shallit. *Algorithmic number theory: efficient algorithms*. Vol. 1. MIT press, 1996 (cit. on p. 57).

[Buc88] J. Buchmann. “A subexponential algorithm for the determination of class groups and regulators of algebraic number fields.” In: *Séminaire de Théorie des Nombres, Paris 1989* (1988), pp. 28–41 (cit. on pp. 204–207).

[Cas12] J. Cassels. *An introduction to the geometry of numbers*. Classics in Mathematics. Springer Berlin Heidelberg, 2012 (cit. on p. 282).

[CDW17] R. Cramer, L. Ducas, and B. Wesolowski. “Short stickelberger class relations and application to Ideal-SVP.” In: *EUROCRYPT*. Springer. 2017, pp. 324–348 (cit. on pp. 86, 89, 168–170).

[CF10] J. Cassels and A. Fröhlich. *Algebraic number theory: proceedings of an instructional conference organized by the london mathematical society (a nato advanced study institute) with the support*

of the international mathematical union. London Mathematical Society, 2010 (cit. on p. 240).

[CG05] H. Cohen and G. Gras. *Class field theory: from theory to practice.* Springer Monographs in Mathematics. Springer Berlin Heidelberg, 2005 (cit. on p. 240).

[Chi08] N. Childress. *Class field theory.* Universitext. Springer New York, 2008 (cit. on pp. 63, 240).

[Coh93] H. Cohen. *A course in computational algebraic number theory.* Springer, 1993 (cit. on p. 239).

[Coh99] H. Cohen. *Advanced topics in computational number theory.* Graduate Texts in Mathematics. Springer New York, 1999 (cit. on pp. 195, 231, 243, 248, 249).

[Cra+16] R. Cramer, L. Ducas, C. Peikert, and O. Regev. “Recovering short generators of principal ideals in cyclotomic rings.” In: *EUROCRYPT*. Springer. 2016, pp. 559–585 (cit. on pp. 36, 86, 89, 168, 170, 188).

[CS08] H. Cohen and P. Stevenhagen. *Computational class field theory.* 2008 (cit. on pp. 204, 231, 240, 260).

[CS10] P. C. Caranay and R. Scheidler. “An efficient seventh power residue symbol algorithm.” In: *International Journal of Number Theory* 6.08 (2010), pp. 1831–1853 (cit. on p. 236).

[CSV12] X. Chang, D. Stehlé, and G. Villard. “Perturbation analysis of the QR factor R in the context of LLL lattice basis reduction.” In: *Math. comput.* 81.279 (2012), pp. 1487–1511 (cit. on pp. 95, 125, 126).

[CT06] T. M. Cover and J. A. Thomas. *Elements of information theory 2nd edition (wiley series in telecommunications and signal processing).* Wiley-Interscience, July 2006 (cit. on p. 185).

[DE16] A. Deitmar and S. Echterhoff. *Principles of harmonic analysis.* 2nd. Springer Publishing Company, Incorporated, 2016 (cit. on pp. 42, 44, 46, 49, 66, 155, 158).

[DF05] I. B. Damgård and G. S. Frandsen. “Efficient algorithms for the gcd and cubic residuosity in the ring of Eisenstein integers.” In: *Journal of Symbolic Computation* 39.6 (2005), pp. 643–652 (cit. on p. 236).

[Dob79] E. Dobrowolski. “On a question of Lehmer and the number of irreducible factors of a polynomial.” In: *Acta Arithmetica* 34.4 (1979), pp. 391–401 (cit. on pp. 187, 188).

[DPW19] L. Ducas, M. Plançon, and B. Wesolowski. “On the shortness of vectors to be found by the Ideal-SVP quantum algorithm.” In: *CRYPTO*. Springer. 2019, pp. 322–351 (cit. on pp. 86, 169).

[Dus98] P. Dusart. “Autour de la fonction qui compte le nombre de nombres premiers.” PhD thesis. l’Université de Limoges: l’Université de Limoges, May 1998 (cit. on pp. 284, 291).

[Eis+14] K. Eisenträger, S. Hallgren, A. Kitaev, and F. Song. “A quantum algorithm for computing the unit group of an arbitrary degree number field.” In: *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*. ACM. 2014, pp. 293–302 (cit. on pp. 34, 35, 81–83, 85–89, 91, 94, 95, 166, 170).

[Gen09] C. Gentry. “A fully homomorphic encryption scheme.” crypto.stanford.edu/craig. PhD thesis. Stanford University, 2009 (cit. on p. 170).

[Gen10] C. Gentry. “Toward basing fully homomorphic encryption on worst-case hardness.” In: *Crypto*. 2010, pp. 116–137 (cit. on pp. 170, 171).

[GKP94] R. L. Graham, D. E. Knuth, and O. Patashnik. *Concrete mathematics: a foundation for computer science*. 2nd. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 1994 (cit. on p. 119).

[GM15] L. Grenié and G. Molteni. “Explicit versions of the prime ideal theorem for Dedekind zeta functions under GRH.” In: *Mathematics of Computation* 85.298 (Oct. 2015), pp. 889–906 (cit. on pp. 57, 58).

[GM84] S. Goldwasser and S. Micali. “Probabilistic encryption.” In: *Journal of Computer and System Sciences* 28.2 (1984), pp. 270–299 (cit. on p. 236).

[GP01] J. von zur Gathen and D. Panario. “Factoring polynomials over finite fields: a survey.” In: *Journal of Symbolic Computation* 31.1 (2001), pp. 3–17 (cit. on p. 59).

[GPV08] C. Gentry, C. Peikert, and V. Vaikuntanathan. “Trapdoors for hard lattices and new cryptographic constructions.” In: *STOC*. 2008, pp. 197–206 (cit. on pp. 176, 195, 196).

[GR02] L. Grover and T. Rudolph. “Creating superpositions that correspond to efficiently integrable probability distributions.” In: *arXiv preprint quant-ph/0208112* (2002) (cit. on pp. 88, 95, 300).

[Gut09] A. Gut. *An intermediate course in probability*. Springer Texts in Statistics. Springer New York, 2009 (cit. on p. 309).

[Hal05] S. Hallgren. “Fast quantum algorithms for computing the unit group and class group of a number field.” In: *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing*. ACM. 2005, pp. 468–474 (cit. on p. 85).

[Hal07] S. Hallgren. “Polynomial-time quantum algorithms for Pell’s equation and the principal ideal problem.” In: *Journal of the ACM (JACM)* 54.1 (2007), p. 4 (cit. on p. 85).

[Hei04] J. Heinonen. *Lectures on Lipschitz analysis*. 2004 (cit. on p. 80).

[HH00] L. Hales and S. Hallgren. “An improved quantum Fourier transform algorithm and applications.” In: *Proceedings 41st Annual Symposium on Foundations of Computer Science*. Nov. 2000, pp. 515–525 (cit. on pp. 87, 116).

[HM89] J. L. Hafner and K. S. McCurley. “A rigorous subexponential algorithm for computation of class groups.” In: *Journal of the American Mathematical Society* 2.4 (1989), pp. 837–850 (cit. on p. 207).

[HM91] J. L. Hafner and K. S. McCurley. “Asymptotically fast triangularization of matrices over rings.” In: *Siam journal on computing* 20.6 (1991), pp. 1068–1083 (cit. on p. 239).

[IKS04] H. Iwaniec, E. Kowalski, and A. M. Society. *Analytic number theory*. American Mathematical Society, 2004 (cit. on pp. 56, 144, 145, 152).

[JMV09] D. Jao, S. D. Miller, and R. Venkatesan. “Expander graphs based on GRH with an application to elliptic curve cryptography.” In: *Journal of number theory* (2009) (cit. on pp. 133, 170).

[JW15] D. Jetchev and B. Wesolowski. “On graphs of isogenies of principally polarizable Abelian surfaces and the discrete logarithm problem.” In: *Corr abs/1506.00522* (2015) (cit. on pp. 133, 162, 170).

[Kes91] V. Kessler. “On the minimum of the unit lattice.” In: *Séminaire de théorie des nombres de bordeaux* 3.2 (1991), pp. 377–380 (cit. on pp. 187, 188).

[Kle00] P. N. Klein. “Finding the closest lattice vector when it’s unusually close.” In: *Soda*. 2000, pp. 937–941 (cit. on pp. 176, 195).

[Koc97] H. Koch. *Algebraic number theory*. Ed. by A. Parshin and I. Shafarevich. Algebraic Number Theory v. 62. Springer Berlin Heidelberg, 1997 (cit. on p. 240).

[Kup05] G. Kuperberg. “A subexponential-time quantum algorithm for the dihedral hidden subgroup problem.” In: *SIAM Journal on Computing* 35.1 (2005), pp. 170–188 (cit. on p. 84).

[KW08] A. Kitaev and W. A. Webb. “Wavefunction preparation and resampling using a quantum computer.” In: *arXiv preprint arXiv:0801.0342* (2008) (cit. on pp. 82, 88, 95, 300, 302).

[Lan12] S. Lang. *Algebraic number theory*. Graduate Texts in Mathematics. Springer New York, 2012 (cit. on pp. 55, 145, 148).

[Lee+19] C. Lee, A. Pellet-Mary, D. Stehlé, and A. Wallet. “An LLL algorithm for module lattices.” In: *ASIACRYPT*. Springer. 2019, pp. 59–90 (cit. on pp. 133, 166, 168).

[Lem00] F. Lemmermeyer. *Reciprocity laws: from Euler to Eisenstein*. Springer Monographs in Mathematics. Springer Berlin Heidelberg, 2000 (cit. on pp. 236, 240).

[Len95] H. W. Lenstra Jr. “Computing Jacobi symbols in algebraic number fields.” In: *Nieuw Archief voor Wiskunde* 13 (1995), pp. 421–426 (cit. on pp. 235, 236, 242).

[LL+93] A. K. Lenstra, H. W. Lenstra Jr., et al. *The development of the number field sieve*. Vol. 1554. Springer Science & Business Media, 1993 (cit. on p. 204).

[Lou00] S. Louboutin. “Explicit bounds for residues of Dedekind zeta functions, values of L-functions at $s=1$, and relative class numbers.” In: *Journal of Number Theory* (2000) (cit. on p. 65).

[ME98] M. Mosca and A. Ekert. “The hidden subgroup problem and eigenvalue estimation on a quantum computer.” In: *NASA International Conference on Quantum Computing and Quantum Communications*. Springer. 1998, pp. 174–188 (cit. on p. 84).

[Mic] D. Micciancio. *Lecture notes on lattice algorithms and applications*. Available at <http://cseweb.ucsd.edu/~daniele/classes.html>, last accessed 17 Oct 2014 (cit. on p. 76).

[Mil15] J. C. Miller. “Real cyclotomic fields of prime conductor and their class numbers.” In: *Math. comp.* 84.295 (2015), pp. 2459–2469 (cit. on p. 55).

[Min67] H. Minkowski. *Gesammelte abhandlungen*. Chelsea, New York, 1967 (cit. on p. 53).

[MR07] D. Micciancio and O. Regev. “Worst-case to average-case reductions based on Gaussian measures.” In: *Siam j. comput.* 37.1 (Apr. 2007), pp. 267–302 (cit. on pp. 77, 78, 86, 178, 188, 200, 309).

[MS18] S. D. Miller and N. Stephens-Davidowitz. “Generalizations of Banaszczyk’s transference theorems and tail bound.” In: *arXiv preprint arXiv:1802.05708* (2018) (cit. on p. 77).

[MV73] H. L. Montgomery and R. C. Vaughan. “The large sieve.” In: *Mathematika* 20.2 (1973), pp. 119–134 (cit. on p. 289).

[Nar04] W. Narkiewicz. *Elementary and analytic theory of algebraic numbers*. Springer Monographs in Mathematics. Springer Berlin Heidelberg, 2004 (cit. on p. 285).

[Nat17] National Institute of Standards and Technology. *Post-quantum cryptography standardization*. 2017 (cit. on p. 84).

[NC11] M. A. Nielsen and I. L. Chuang. *Quantum computation and quantum information: 10th anniversary edition*. 10th. New York, NY, USA: Cambridge University Press, 2011 (cit. on pp. 44, 301, 304, 305).

[Neu85] J. Neukirch. *Class field theory*. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 1985 (cit. on pp. 56, 237, 240, 260, 262).

[NS13] J. Neukirch and N. Schappacher. *Algebraic number theory*. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 2013 (cit. on pp. 1, 4, 53, 56, 65, 149, 152, 240).

[NS16] A. Neumaier and D. Stehlé. “Faster LLL-type reduction of lattice bases.” In: *Proceedings of the acm on international symposium on symbolic and algebraic computation*. 2016, pp. 373–380 (cit. on pp. 95, 231).

[Ove14] M. Overholt. *A course in analytic number theory*. Graduate Studies in Mathematics. American Mathematical Society, 2014 (cit. on p. 209).

[PAR19] *PARI/GP version 2.11.2*. Available at <http://pari.math.u-bordeaux.fr/>. The PARI Group. Univ. Bordeaux, 2019 (cit. on p. 204).

[Pei16] C. Peikert. “A decade of lattice cryptography.” In: *Foundations and Trends in Theoretical Computer Science* 10.4 (2016), pp. 283–424 (cit. on p. 1).

[PHS19] A. Pellet-Mary, G. Hanrot, and D. Stehlé. “Approx-SVP in ideal lattices with pre-processing.” In: *EUROCRYPT*. Springer. 2019, pp. 685–716 (cit. on pp. 86, 133, 166, 168, 169).

[RA08] M. Reiter and S. Arthur. *Fourier transform & Sobolev spaces (lecture notes)*. Available at: https://www.mat.univie.ac.at/~stein/teaching/SoSem08/sobolev_fourier.pdf. 2008 (cit. on p. 80).

[Rab89] S. Rabinowitz. “The volume of an n-simplex with many equal edges.” In: *Missouri Journal of Mathematical Sciences* 1 (Jan. 1989) (cit. on p. 279).

[Reg04a] O. Regev. *Lecture notes in ‘lattices in computer science’*. Available at https://cims.nyu.edu/~regev/teaching/lattices_fall_2004/. Nov. 2004 (cit. on p. 122).

[Reg04b] O. Regev. “Quantum computation and lattice problems.” In: *SIAM Journal on Computing* 33.3 (2004), pp. 738–760 (cit. on p. 84).

[Reg09] O. Regev. “On lattices, learning with errors, random linear codes, and cryptography.” In: *J. ACM* 56.6 (2009). Preliminary version in STOC 2005, pp. 1–40 (cit. on p. 170).

[Sch08] R. Schoof. “Computing Arakelov class groups.” In: *Algorithmic Number Theory: Lattices, Number Fields, Curves and Cryptography*. Cambridge University Press. 2008, pp. 447–495 (cit. on pp. 1, 74, 132, 166, 207).

[Sch98] R. Scheidler. “A public-key cryptosystem using purely cubic fields.” In: *Journal of Cryptology* 11.2 (1998), pp. 109–124 (cit. on p. 236).

[Ser77] J.-P. Serre. *Linear representations of finite groups*. Vol. 42. Graduate texts in mathematics. Springer, 1977, pp. I–X, 1–170 (cit. on p. 155).

[Sey87] M. Seysen. “A probabilistic factorization algorithm with quadratic forms of negative discriminant.” In: *Mathematics of Computation* 48.178 (1987), pp. 757–780 (cit. on p. 207).

[Sho94] P. W. Shor. “Algorithms for quantum computation: discrete logarithms and factoring.” In: *Proceedings 35th Annual Symposium on Foundations of Computer Science*. IEEE. 1994, pp. 124–134 (cit. on pp. 84, 236).

[Sho95] V. Shoup. “A new polynomial factorization algorithm and its implementation.” In: *Journal of symbolic computation* 20.4 (1995), pp. 363–397 (cit. on p. 59).

[SL96] A. Storjohann and G. Labahn. “Asymptotically fast computation of Hermite normal forms of integer matrices.” In: *Proceedings of the 1996 international symposium on symbolic and algebraic computation*. ISSAC ’96. Zurich, Switzerland: Association for Computing Machinery, 1996, pp. 259–266 (cit. on pp. 231, 239).

[Son13] F. Song. “Quantum computing: a cryptographic perspective.” Available at: https://etda.libraries.psu.edu/files/final_submissions/8820. PhD thesis. The Pennsylvania State University, 2013 (cit. on p. 89).

[Squ97] D. Squirrel. *An algorithm for the power residue symbol*. 1997 (cit. on pp. 235, 236, 242, 246, 249).

[SV05] A. Schmidt and U. Vollmer. “Polynomial time quantum algorithm for the computation of the unit group of a number field.” In: *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing*. ACM. 2005, pp. 475–480 (cit. on p. 85).

[SW95] R. Scheidler and H. C. Williams. “A public-key cryptosystem utilizing cyclotomic fields.” In: *Designs, Codes and Cryptography* 6.2 (1995), pp. 117–131 (cit. on p. 236).

[Vil85] A. Villani. “Another note on the inclusion $L^p(\mu) \subset L^q(\mu)$.” In: *The American Mathematical Monthly* 92.7 (1985), pp. 485–487 (cit. on p. 80).

[Was12] L. C. Washington. *Introduction to cyclotomic fields*. Vol. 83. Springer Science & Business Media, 2012 (cit. on pp. 187, 188).

Bibliography

[Wei02] A. Weilert. “Fast computation of the biquadratic residue symbol.” In: *Journal of Number Theory* 96.1 (2002), pp. 133–151 (cit. on p. [236](#)).

[Wer07] D. Werner. *Funktionalanalysis*. Springer-Lehrbuch. Springer Berlin Heidelberg, 2007 (cit. on p. [80](#)).

[Wes18] B. P. Wesolowski. “Arithmetic and geometric structures in cryptography.” PhD thesis. École Polytechnique Fédérale de Lausanne, Nov. 2018 (cit. on p. [144](#)).

[Wil85] H. C. Williams. “An M^3 public-key encryption scheme.” In: *Conference on the Theory and Application of Cryptographic Techniques*. Springer. 1985, pp. 358–368 (cit. on p. [236](#)).

[Yud76] V. A. Yudin. “The multidimensional Jackson theorem.” In: *Mathematical notes of the Academy of Sciences of the USSR* 20.3 (Sept. 1976), pp. 801–804 (cit. on pp. [52](#), [294](#), [295](#)).

Appendix

Appendix A.

Appendix

A.1. Number-theoretic Computations

Lemma A.1. *The volume of the simplex $S_\alpha = \{x \in \text{Log } K_{\mathbb{R}} \mid x_\sigma \leq \alpha, \sum_\sigma x_\sigma = 0\}$ for some $\alpha > 0$ is given by*

$$\text{Vol}(S_\alpha) = \frac{(n\alpha)^r \cdot \sqrt{n}}{\sqrt{2}^{n_{\mathbb{C}}} \cdot r!},$$

where $r = n_{\mathbb{R}} + n_{\mathbb{C}} - 1$.

Proof. Define $S'_\alpha = \{x \in \mathbb{R}^{r+1} \mid \sum_\nu x_\nu = 0, x_\nu \leq \alpha \text{ for real places } \nu, x_\nu \leq 2\alpha \text{ for complex } \nu\}$. The map

$$A : \mathbb{R}^{r+1} \rightarrow \text{Log } K_{\mathbb{R}}, e_\nu \mapsto \begin{cases} e_{\sigma_\nu} & \text{when } \nu \text{ is real} \\ \frac{1}{2}(e_{\sigma_\nu} + e_{\bar{\sigma}_\nu}) & \text{when } \nu \text{ is complex} \end{cases}$$

sends S'_α to S_α bijectively. By applying on $S'_\alpha \subseteq \mathbb{R}^{r+1}$ the translation $y_\sigma = \alpha - x_\sigma$ or $y_\sigma = 2\alpha - x_\sigma$ depending on whether σ is real or complex, one can see that it is a regular r -simplex with edge length $\sqrt{2} \cdot n\alpha$. Therefore, the volume of S'_α equals $\frac{(n\alpha)^r \sqrt{r+1}}{r!}$ [Rab89]. In order to compute the volume of S_α , we need to estimate how the linear map A scales the subspace $\{x \in \mathbb{R}^{r+1} \mid \sum_\nu x_\nu = 0\}$. Therefore, we choose the basis $B = (e_1 - e_{r+1}, \dots, e_r - e_{r+1})$, and compute the scaling factor by means of taking the square root

of the determinant of $(AB)^T AB$ and dividing it by the square root of the determinant of $B^T B$, i.e.,

$$\text{Vol}(S_\alpha) = \frac{\sqrt{\det(B^T A^T AB)}}{\sqrt{\det(B^T B)}} \text{Vol}(S'_\alpha).$$

By the Weinstein–Aronszajn identity, we obtain that $\det(B^T B) = \det(I + \mathbf{1} \cdot \mathbf{1}^T) = n_{\mathbb{R}} + n_{\mathbb{C}} = \mathfrak{r} + 1$, where $\mathbf{1}$ is the all-one column vector of dimension $\mathfrak{r} = n_{\mathbb{R}} + n_{\mathbb{C}} - 1$. Note that $A^T A = \text{diag}(1, \dots, 1, 1/2, \dots, 1/2)$, where the 1 is repeated $n_{\mathbb{R}}$ times and the $1/2$ is repeated $n_{\mathbb{C}}$ times. Therefore, $B^T A^T AB = J + \frac{1}{2} \mathbf{1} \cdot \mathbf{1}^T$, where $J = \text{diag}(\underbrace{1, \dots, 1}_{n_{\mathbb{R}}}, \underbrace{1/2, \dots, 1/2}_{n_{\mathbb{C}}-1})$. Again using the Weinstein–Aronszajn identity, we obtain

$$\begin{aligned} \det(B^T A^T AB) &= \det(J + 1/2 \cdot \mathbf{1} \cdot \mathbf{1}^T) = \det(J)(1 + 1/2 \cdot \mathbf{1}^T J^{-1} \mathbf{1}) \\ &= 2^{-n_{\mathbb{C}}+1} (1 + 1/2(n_{\mathbb{R}} + 2n_{\mathbb{C}} - 2)) = 2^{-n_{\mathbb{C}}} \cdot n \end{aligned}$$

So, we conclude the argument by spelling out all formula's:

$$\text{Vol}(S_\alpha) = \frac{2^{-n_{\mathbb{C}}} \sqrt{n}}{\sqrt{\mathfrak{r}+1}} \text{Vol}(S'_\alpha) = \frac{2^{-n_{\mathbb{C}}} \sqrt{n}}{\sqrt{\mathfrak{r}+1}} \cdot \frac{(n\alpha)^{\mathfrak{r}} \sqrt{\mathfrak{r}+1}}{\mathfrak{r}!} = \frac{(n\alpha)^{\mathfrak{r}} \cdot \sqrt{n}}{\sqrt{2^{n_{\mathbb{C}}}} \cdot \mathfrak{r}!}$$

□

Lemma A.2. *Let $\text{Log } \mathcal{O}_K^\times \subseteq H \subseteq \log K_{\mathbb{R}}$ be the logarithmic unit lattice. Then the covolume of this lattice in H equals $\sqrt{n} \cdot 2^{-n_{\mathbb{C}}/2} \cdot R$.*

Proof. In the literature, often one uses the embedding $\text{Log}' \mathcal{O}_K^\times \subseteq H' \subseteq \mathbb{R}^{n_{\mathbb{R}}+n_{\mathbb{C}}}$, where $(\text{Log}'(\eta))_\sigma$ equals $\log |\sigma(\eta)|$ or $2 \log |\sigma(\eta)|$, depending on whether σ is real or complex. The space $H' = \{x \in \mathbb{R}^{n_{\mathbb{R}}+n_{\mathbb{C}}} \mid \sum_j x_j = 0\}$ is the equivalent hyperplane. It is evident that the linear map

$$A : \mathbb{R}^{\mathfrak{r}+1} \rightarrow \text{Log } K_{\mathbb{R}}, e_\nu \mapsto \begin{cases} e_{\sigma_\nu} & \text{when } \nu \text{ is real} \\ \frac{1}{2}(e_{\sigma_\nu} + e_{\bar{\sigma}_\nu}) & \text{when } \nu \text{ is complex} \end{cases}$$

maps $\text{Log}' \mathcal{O}_K^\times \subseteq H'$ to $\text{Log } \mathcal{O}_K^\times \subseteq H$.

Let \underline{U} be a basis of $\text{Log}' \mathcal{O}_K^\times$, and denote U by the same basis, but the last row removed; the determinant of U is called the regulator R of the number field K . Define $B : \mathbb{R}^r \rightarrow \mathbb{R}^{r+1}$, $e_j \mapsto e_j - e_{n_{\mathbb{R}}+n_{\mathbb{C}}}$. By the fact that for any element in $\text{Log}' \mathcal{O}_K^\times$ holds that the sum of the entries equals zero, we have $BU = \underline{U}$. As A maps $\text{Log}' \mathcal{O}_K^\times$ to $\text{Log} \mathcal{O}_K^\times$, we obtain that ABU is a basis of $\text{Log} \mathcal{O}_K^\times$. The covolume of this basis equals $\sqrt{\det(B^T A^T AB)} \det(U) = \sqrt{\det(B^T A^T AB)} R = \sqrt{n} 2^{-n_{\mathbb{C}}/2} R$.

The last equality is proven by the computation of $\det(B^T A^T AB)$ below. Note that $A^T A = \text{diag}(1, \dots, 1, 1/2, \dots, 1/2)$, where the 1 is repeated $n_{\mathbb{R}}$ times and the $1/2$ is repeated $n_{\mathbb{C}}$ times. Therefore, $B^T A^T AB = J + \frac{1}{2} \mathbf{1} \cdot \mathbf{1}^T$, where

$$J = \text{diag}(\underbrace{1, \dots, 1}_{n_{\mathbb{R}}}, \underbrace{1/2, \dots, 1/2}_{n_{\mathbb{C}}-1}).$$

and $\mathbf{1}$ is the all-one vector of dimension r . Using the Weinstein-Aronszajn identity, we obtain

$$\begin{aligned} \det(B^T A^T AB) &= \det(J + 1/2 \cdot \mathbf{1} \cdot \mathbf{1}^T) = \det(J)(1 + 1/2 \cdot \mathbf{1}^T J^{-1} \mathbf{1}) \\ &= 2^{-n_{\mathbb{C}}+1} (1 + \frac{1}{2} (n_{\mathbb{R}} + 2n_{\mathbb{C}} - 2)) = 2^{-n_{\mathbb{C}}} \cdot n \end{aligned}$$

□

Lemma A.3. *Let $H \subseteq \text{Log}(K_{\mathbb{R}})$ be the hyperplane orthogonal to the all-one vector, and let $\rho_s^{(n)}$ be the Gaussian function. Then*

$$\int_{x \in H} s^{-r} \rho_s^{(n)}(x) dx = 1$$

Proof. Use the matrices A and B from the previous lemma to apply integration by substitution, observing that $H = AB\mathbb{R}^r$.

$$\begin{aligned} \int_{x \in AB\mathbb{R}^r} s^{-r} \rho_s^{(n)}(x) dx &= \sqrt{\det(B^T A^T AB)} \int_{x \in \mathbb{R}^r} s^{-r} \rho_s^{(n)}(ABx) dx \\ &= \sqrt{\det(D^T D)} \int_{x \in \mathbb{R}^r} s^{-r} e^{-\pi x^T D^T D x / s^2} dx = \int_{x \in \mathbb{R}^r} s^{-r} e^{-\pi x^T x / s^2} dx = 1 \end{aligned}$$

Where $D^T D = B^T A^T A B^T$ is the r -dimensional Cholesky decomposition, and the last equality follows then again by integration by substitution. \square

Theorem A.4 (Bhargava,Shankar,Taniguchi,Thorne,Tsimerman,Zhao). *Let K be any number field of degree n and let \mathcal{O}_K be its ring of integers. Let $\mathcal{O}_K \subseteq K_{\mathbb{R}}$ have the structure of a lattice via the Minkowski embedding (see Section 2.3), and denote $\lambda_j^{\infty}(\mathcal{O}_K)$ for the j -th successive minimum with respect to the infinity norm in $K_{\mathbb{R}}$. Then*

$$\lambda_n^{\infty}(\mathcal{O}_K) \leq |\Delta_K|^{1/n}.$$

The following proof is a copy of [Bha+20, Thm. 3.1], with the difference that it is applied to the infinity norm and has explicit constants everywhere.

Proof. Let $\alpha_j \in \mathcal{O}_K$ attain the successive minima for the infinity norm $\lambda_j^{\infty}(\mathcal{O}_K)$ for $j \in \{1, \dots, n\}$, with $\alpha_1 = 1$. For any element $\beta \in \mathcal{O}_K$, we write $\beta = \sum_{j=1}^n [\beta]_j \alpha_j$, i.e., $[\beta]_j$ are the coordinates of β with respect to $(\alpha_1, \dots, \alpha_n)$.

For $2 \leq k, \ell \leq n-1$ consider the $(n-2) \times (n-2)$ -matrix $C = ([\alpha_k \alpha_{\ell}]_n)$, i.e., the matrix consisting of the coordinates of $\alpha_k \alpha_{\ell}$ with respect to α_n . We will show at the end of this proof that this is a non-degenerate matrix, implying that there are no zero rows or columns. In other words, there exists a permutation $\pi : \{2, \dots, n-1\} \rightarrow \{2, \dots, n-1\}$ such that $[\alpha_k \alpha_{\pi(k)}]_n \neq 0$ for all $k \in \{2, \dots, n-1\}$.

So, the product $\alpha_k \alpha_{\pi(k)} \in \mathcal{O}_K$ extends $\{\alpha_1, \dots, \alpha_{n-1}\}$ to a n -dimensional lattice; therefore we have $\|\alpha_k\|_{\infty} \|\alpha_{\pi(k)}\|_{\infty} \geq \|\alpha_k \alpha_{\pi(k)}\|_{\infty} \geq \lambda_n^{\infty}(\mathcal{O}_K)$. Taking products over all $k \in \{2, \dots, n-1\}$ we obtain

$$\prod_{k=2}^{n-1} \|\alpha_k\|_{\infty}^2 = \prod_{k=2}^{n-1} \|\alpha_k\|_{\infty} \|\alpha_{\pi(k)}\|_{\infty} \geq (\lambda_n^{\infty}(\mathcal{O}_K))^{n-2}.$$

Multiplying above equation by $\|\alpha_1\|_{\infty}^2 = 1$ and $\|\alpha_n\|_{\infty}^2 = \lambda_n^{\infty}(\mathcal{O}_K)^2$, and using Minkowski's second inequality [Cas12, Ch. VIII] $\prod_{k=1}^n \lambda_k^{\infty}(\Lambda) \leq \det(\Lambda)$,

we obtain

$$|\Delta_K| \geq \prod_{k=1}^n \|\alpha_k\|_\infty^2 \geq (\lambda_n^\infty(\mathcal{O}_K))^n.$$

It remains to prove that $C = ([\alpha_k \alpha_\ell]_n)$ is non-degenerate. Suppose it is not, and there exists d_ℓ for $\ell \in \{2, \dots, n-1\}$ (not all zero) such that

$$\left[\sum_{\ell=2}^{n-1} d_\ell \alpha_k \alpha_\ell \right]_n = \sum_{\ell=2}^{n-1} d_\ell [\alpha_k \alpha_\ell]_n = 0 \text{ for all } k \in \{2, \dots, n-1\}$$

Writing $\beta = \sum_{\ell=2}^{n-1} d_\ell \alpha_\ell$, this means that $\alpha_k \beta$ lies in the span of the elements $(\alpha_1, \dots, \alpha_{n-1})$. In other words, $L = \mathbb{Q}\alpha_1 + \dots + \mathbb{Q}\alpha_{n-1}$ is $\mathbb{Q}(\beta)$ -invariant, i.e., a $\mathbb{Q}(\beta)$ -vector (strict) subspace of K . That is, $\dim_{\mathbb{Q}(\beta)}(L) \leq \dim_{\mathbb{Q}(\beta)}(K) - 1$. But then

$$\begin{aligned} n-1 &= \dim_{\mathbb{Q}}(L) = \dim_{\mathbb{Q}(\beta)}(L) \cdot [\mathbb{Q}(\beta) : \mathbb{Q}] \\ &\leq (\dim_{\mathbb{Q}(\beta)}(K) - 1) \cdot [\mathbb{Q}(\beta) : \mathbb{Q}] = n - [\mathbb{Q}(\beta) : \mathbb{Q}], \end{aligned}$$

yielding $[\mathbb{Q}(\beta) : \mathbb{Q}] = 1$, i.e., $\beta \in \mathbb{Q}$, which is impossible by the fact that $\beta = \sum_{\ell=2}^{n-1} d_\ell \alpha_\ell$ is assumed to be non-zero and has no $\alpha_1 = 1$ part.

We conclude that C is non-degenerate, which finishes the proof. □

A.2. Bound on the Residue of the Zeta Function for Cyclotomic Fields

In the proof of Lemma 7.20, we used that for the cyclotomic field $K = \mathbb{Q}(\zeta_m)$, the residue ρ_K of the zeta function for cyclotomic fields is in $O(m^4)$. This section is dedicated to the proof of this fact.

Theorem A.5 (ERH). *Let $K = \mathbb{Q}(\zeta_m)$ with $m \geq 3$. Then, assuming the Riemann Hypothesis for L-functions $L(\chi, s)$ for all Dirichlet characters modulo m , we have*

$$\rho_K \leq e^{15} \cdot m^4 = O(m^4).$$

Proof. The proof extends to the rest of this section, through the following steps.

(Appendix A.2.1) Writing $\log(\rho_K) = R_K + M_K$

We first split the computation of ρ_K into two parts, a *ramified part* R_K and a *main part* M_K . This ramified part occurs because the characters $\chi \in \hat{G} \setminus \{1\}$ for $G = \text{Gal}(K/\mathbb{Q})$ are defined modulo their conductor $f_\chi \mid m$. For computations it is simpler to consider characters modulo m instead, denoted, $\chi|_m$. The ramified term pops up as a correction factor, just being the sum of $L(\chi, 1) - L(\chi|_m, 1)$ for the non-trivial characters χ .

(Appendix A.2.2) Bounding the ramified term $R_K \leq 2 \log(m)$

By elementary methods one can show that $R_K \leq 2 \log(m)$ (see Proposition A.9).

(Appendix A.2.3) Splitting $M_K = M_K^{(w)} + \lim_{x \rightarrow \infty} (M_K^{(x)} - M_K^{(w)})$.

The main part $M_K = \sum_q \frac{a_q}{q}$ can be seen as a sum where q ranges over all prime powers. By defining the partial sum $M_K^{(w)} = \sum_{q < w} \frac{a_q}{q}$ one obtains an ‘initial’ part $M_K^{(w)}$ and a ‘tail part’ $\lim_{x \rightarrow \infty} (M_K^{(x)} - M_K^{(w)})$ of M_K .

(Appendix A.2.4) The initial part $M_K^{(w)} \leq 2 \log(m) + 11$ for $w = \max(e^{5/4 \cdot m}, 10^{10})$.

By applying partial summation to the Brun-Titchmarsh bound (see Lemma A.13) one obtains the bound $M_K^{(w)} \leq 2 \log \log w + 7$. It is easy to show that for $w = \max(e^{5/4 \cdot m}, 10^{10})$ holds $2 \log \log w + 7 \leq 2 \log(m) + 11$.

(Appendix A.2.5) The tail part $\lim_{x \rightarrow \infty} (M_K^{(x)} - M_K^{(w)}) \leq 4$ for $w = \max(e^{5/4 \cdot m}, 10^{10})$.

This bound, proven in Proposition A.17, assumes the Riemann Hypothesis for L -functions for Dirichlet characters modulo m , and follows from an explicit result of Dusart [Dus98].

Combining the bounds yields $\log(\rho_K) \leq 4 \log(m) + 15$.

We have the following bound, of which taking the exponent yields the final claim.

$$\log \rho_K \leq R_K + M_K^{(w)} + \lim_{x \rightarrow \infty} (M_K^{(x)} - M_K^{(w)}) \leq 2 \log(m) + (2 \log(m) + 11) + 4.$$

□

A.2.1. Splitting $\log(\rho_K) = R_K + M_K$ into a Ramified Term and a Main Term

Notation A.6. In the following, every Dirichlet character χ is assumed to be primitive, i.e., defined modulo its conductor f_χ . If we, instead, want to consider a Dirichlet character modulo a larger modulus m (with $f_\chi \mid m$), we write $\chi|_m$ (and we have $\chi|_m(a) = 0$ whenever $\gcd(a, m) > 1$). We denote by $\mathbf{1}$ the trivial character that has value one everywhere.

Lemma A.7. Let $K = \mathbb{Q}(\zeta_m)$ be a cyclotomic field extension with Galois group $G \simeq (\mathbb{Z}/m\mathbb{Z})^*$ and consider all characters \hat{G} as Dirichlet characters. Then we have $\log(\rho_K) = R_K + M_K$, where

$$R_K = - \sum_{\substack{\chi \in \hat{G} \setminus \mathbf{1} \\ p \mid m \\ p \nmid f_\chi}} \sum_{p \mid m} \log(1 - \chi(p)/p) \text{ and } M_K = \sum_{\chi \in \hat{G} \setminus \mathbf{1}} \log L(\chi|_m, 1)$$

Proof. We have the following formula for the logarithm of the residue ρ_K , by considering the quotient of the Dedekind zeta function and the Riemann zeta function [Nar04, Thm. 8.6].

$$\log(\rho_K) = \sum_{\chi \in \hat{G} \setminus \mathbf{1}} \log L(\chi, 1)$$

Concentrating on a fixed $\chi \in \hat{G} \setminus \mathbf{1}$, and applying the Euler product formula, we obtain

$$\begin{aligned} \log L(\chi, 1) &= - \sum_{p \nmid f_\chi} \log(1 - \chi(p)/p) \\ &= - \sum_{\substack{p \mid m \\ p \nmid f_\chi}} \log(1 - \chi(p)/p) - \sum_{\substack{p \mid m \\ p \mid f_\chi}} \log(1 - \chi(p)/p) \\ &= \log L(1, \chi|_m) - \sum_{\substack{p \mid m \\ p \nmid f_\chi}} \log(1 - \chi(p)/p). \end{aligned}$$

Summing over all non-trivial $\chi \in \hat{G}$ yields

$$\log(\rho_K) = - \sum_{\substack{\chi \in \hat{G} \setminus \mathbf{1} \\ p \nmid f_\chi}} \sum_{p|m} \log(1 - \chi(p)/p) + \sum_{\substack{\chi \in \hat{G} \setminus \mathbf{1}}} \log L(1, \chi|_m) = R_K + M_K.$$

□

We call the terms R_K and M_K the *ramified term* and the *main term* respectively.

A.2.2. Estimating the Ramified Term

Lemma A.8. *For any prime-power cyclotomic number field $K = \mathbb{Q}(\zeta_{p^k})$, the ramified term R_K equals zero.*

Proof. For a prime-power cyclotomic field $\mathbb{Q}(\zeta_{p^k})$, the conductor of every non-trivial character $\chi \in \hat{G}$ is divisible by p , since $G = \text{Gal}(\mathbb{Q}(\zeta_{p^k})/\mathbb{Q}) \simeq (\mathbb{Z}/p^k\mathbb{Z})^*$. Therefore, $R_K = - \sum_{\chi \in \hat{G} \setminus \mathbf{1}} \sum_{\substack{p|m \\ p \nmid f_\chi}} \log(1 - \chi(p)/p) = 0$. □

Proposition A.9. *For any cyclotomic number field $K = \mathbb{Q}(\zeta_m)$ with $m \geq 3$, we have*

$$R_K \leq 2 \log(m)$$

Proof. Denoting $G \simeq (\mathbb{Z}/m\mathbb{Z})^*$ for the Galois group of K , swapping sums and using the Taylor expansion of the logarithm, we obtain

$$\begin{aligned} R_K &= \sum_{p|m} \sum_{\substack{\chi \in \hat{G} \setminus \mathbf{1} \\ p \nmid f_\chi}} \sum_{j>0} \frac{\chi(p^j)}{jp^j} = \sum_{p|m} \sum_{j>0} \frac{1}{jp^j} \sum_{\substack{\chi \in \hat{G} \setminus \mathbf{1} \\ p \nmid f_\chi}} \chi(p^j) \\ &= \sum_{p|m} \sum_{j>0} \frac{1}{jp^j} \left(-1 + \sum_{\chi \in \hat{G}_p} \chi(p^j) \right). \end{aligned}$$

where $\hat{G}_p = \{\chi \in \hat{G} \mid p \nmid f_\chi\}$. Note that $\hat{G}_p \simeq (\mathbb{Z}/m_p\mathbb{Z})^*$ is isomorphic to the Galois group of $\mathbb{Q}(\zeta_{m_p})$, where m_p is the p -free part of m . By character orthogonality relations, we know that

$$\sum_{\chi \in \hat{G}_p} \chi|_{m_p}(a) = \begin{cases} |\hat{G}_p| = \phi(m_p) & \text{if } a \equiv 1 \pmod{m_p} \\ 0 & \text{otherwise} \end{cases}$$

Since p is coprime with m_p , we know that for any character χ of \hat{G}_p and exponent $j > 0$, it holds that $\chi(p^j) = \chi|_{m_p}(p^j)$. Denoting j_p for the order of p in $(\mathbb{Z}/m_p\mathbb{Z})^*$, we deduce that j_p is the smallest non-zero exponent such satisfying $\sum_{\chi \in \hat{G}_p} \chi(p^{j_p}) = \phi(m_p)$. Moreover, we have $p^{j_p} = 1 + km_p > m_p$. Using these properties, we obtain the following rather crude bound.

$$\begin{aligned} \sum_{p|m} \sum_{j>0} \frac{1}{jp^j} \left(-1 + \sum_{\chi \in \hat{G}_p} \chi(p^j) \right) &\leq \sum_{p|m} \sum_{k>0} \frac{\phi(m_p) - 1}{(kj_p)p^{kj_p}} \\ &\leq - \sum_{p|m} (\phi(m_p) - 1) \log(1 - p^{-j_p}) \\ &\leq \sum_{p|m} \frac{2 \log(2) \cdot (\phi(m_p) - 1)}{p^{j_p}} \\ &\leq 2 \log(2) \cdot \omega(m) \leq 2 \log(m) \end{aligned}$$

The first inequality omits the $p^j \not\equiv 1 \pmod{m_p}$, as they add negative value anyway; the second inequality uses the equation $\sum_{k>0} (p^{-j_p})^k / k = -\log(1 - p^{-j_p})$ after disposing j_p in the denominator. The third inequality uses the fact that $-\log(1 - x) \leq 2 \log(2) \cdot x$ for $x < 1/2$, the fourth inequality uses the fact that $p^{j_p} > m_p$. By Lemma A.8, we may assume, without loss of generality, that m has at least 2 distinct prime divisors, i.e., $\omega(m) > 1$. Then the fifth inequality is just a trivial upper bound on the prime omega function $\omega(m)$, the number of *distinct* prime divisors of m . \square

A.2.3. Splitting the Main Term in an Initial Part and a Tail Part

Notation A.10. For $a \in \mathbb{N}$ with $\gcd(a, m) = 1$, we put

$$S_{a,x} = \sum_{\substack{p \text{ prime}, j > 0, \\ p^j \equiv a \pmod{m} \\ p^j \leq x}} \frac{1}{jp^j}$$

Proposition A.11 (Estimating the main term). Let $K = \mathbb{Q}(\zeta_m)$ be a cyclotomic field. Then

$$M_K = \lim_{x \rightarrow \infty} \left(\phi(m) \cdot S_{1,x} - \sum_{a \in (\mathbb{Z}/m\mathbb{Z})^*} S_{a,x} \right)$$

Proof. We have

$$M_K = \sum_{\chi \in \hat{G} \setminus \{1\}} \log L(\chi|_m, 1) = \sum_{p \nmid m} \sum_j \frac{1}{jp^j} \sum_{\chi \in \hat{G} \setminus \{1\}} \chi|_m(p^j)$$

For numbers a coprime with m we know that $\sum_{\chi \in \hat{G}} \chi|_m(a)$ equals $\phi(m)$ if $a \equiv 1 \pmod{m}$ and 0 otherwise. This yields:

$$M_K = (\phi(m) - 1) \sum_{\substack{p \text{ prime}, j > 0 \\ p^j \equiv 1 \pmod{m}}} \frac{1}{jp^j} - \sum_{\substack{p \text{ prime}, j > 0 \\ p \nmid m, p^j \not\equiv 1 \pmod{m}}} \frac{1}{jp^j}.$$

Writing out the new notation and flipping summands corresponding to $p^j \equiv 1 \pmod{m}$ from the left-hand to the right-hand side yields the result. \square

It will be proven useful to cut the main term into two parts:

$$M_K = M_K^{(w)} + \lim_{x \rightarrow \infty} (M_K^{(x)} - M_K^{(w)}).$$

That is, a finite initial part $M_K^{(w)}$ and a tail part $\lim_{x \rightarrow \infty} (M_K^{(x)} - M_K^{(w)})$. More precisely, for $w > 1$,

Notation A.12.

$$M_K^{(w)} = \phi(m)S_{1,w} - \sum_{b \in (\mathbb{Z}/m\mathbb{Z})^*} S_{b,w}$$

A.2.4. Estimating the Initial Part of the Main Term

Lemma A.13. *For $w \geq m^4$ we have*

$$M_K^{(w)} \leq 2 \log \log w + 7.$$

Proof. By omitting the negative terms in Notation A.12, we obtain

$$M_K^{(w)} \leq \phi(m) S_{1,w} = \phi(m) \sum_{\substack{p \text{ prime}, j > 0 \\ p^j \equiv 1 \pmod{m} \\ p^j \leq w}} \frac{1}{jp^j} \leq 5 + \phi(m) \sum_{\substack{p \text{ prime} \\ p \equiv 1 \pmod{m} \\ p \leq w}} \frac{1}{p}.$$

where the last inequality follows from Lemma A.14

$$\sum_{\substack{p \text{ prime}, j > 1 \\ p^j \equiv 1 \pmod{m}}} \frac{1}{jp^j} \leq 5/m,$$

For a fixed m , we denote by $\pi_1(t)$ the number of primes p with $p \leq t$ that satisfy $p \equiv 1 \pmod{m}$. For $t > m$, we have the Brun-Titchmarsh bound $\pi_1(t) \leq \frac{2t}{\phi(m) \log(t/m)}$ [MV73]. Combining this bound with Abel partial summation, we obtain

$$\begin{aligned} \sum_{\substack{p \text{ prime} \\ p \equiv 1 \pmod{m} \\ p \leq w}} \frac{1}{p} &\leq \frac{1}{m} + \frac{1}{2m} + \sum_{\substack{p \text{ prime} \\ p \equiv 1 \pmod{m} \\ em \leq p \leq w}} \frac{1}{p} \\ &= \frac{1}{m} + \frac{1}{2m} + \frac{\pi_1(w)}{w} - \frac{\pi_1(em)}{em} + \int_{em}^w \frac{2dx}{\phi(m)x \log(x/m)} \\ &\leq \frac{3}{2m} + \frac{1}{\phi(m) \log(w/m)} + 2/\phi(m) \cdot \log \log(w/m) \end{aligned}$$

The first inequality just writes out the first two terms of the sum, the subsequent equality is the Abel summation formula, using the facts that t^{-1} has derivative $-t^{-2}$ and π_1 has the Brun-Titchmarsh bound. The last inequality follows from evaluating the integral, combining the terms and using again the Brun-Titchmarsh bound for $\pi_1(w)$. Concluding, one can deduce that $M_K^{(w)}$ is bounded by $5 + 3/2 + 1/\log(w/m) + 2 \log \log w \leq 7 + 2 \log \log w$. \square

Lemma A.14. *For all $m \geq 2$ holds*

$$\sum_{\substack{p \text{ prime}, j > 1 \\ p^j \equiv 1 \pmod{m} \\ p > m}} \frac{1}{jp^j} \leq \frac{5}{m},$$

Proof. Using the technique from Ankeny and Chowla [AC49, p. 532] we split the sum into a part where $p > m$ and a part where $p < m$.

For $p > m$ we have

$$\sum_{\substack{p \text{ prime}, j > 1 \\ p^j \equiv 1 \pmod{m} \\ p > m}} \frac{1}{jp^j} \leq \sum_{k > m} \frac{1}{k^2} \leq \int_m^\infty 1/x^2 \cdot dx = \frac{1}{m} \quad (\text{A.115})$$

The first inequality follows from the fact that for every fixed prime $p > m$ we have

$$\sum_{j > 1} \frac{1}{jp^j} \leq \frac{1}{2p^2} \left(\sum_{j=0}^{\infty} p^{-j} \right) \leq \frac{1}{2p^2} \cdot \frac{p}{p-1} \leq \frac{1}{p^2}.$$

For $p < m$ we use the fact that $X^k \equiv 1 \pmod{m}$ can have at most k incongruent solutions [AC49, p. 532]. This implies, by considering all numbers $am + 1$ with $a \in \mathbb{Z}$,

$$\sum_{\substack{p \text{ prime}, j > 1 \\ p^j \equiv 1 \pmod{m} \\ p < m}} \frac{1}{jp^j} \leq \sum_{j=2}^{\infty} \frac{1}{j} \left[\sum_{a=A(j)}^{B(j)} \frac{1}{am+1} \right] \leq \sum_{j=2}^{\infty} \frac{1}{(\frac{j^2-j}{2} + 1)m + 1},$$

where $A(j) = \frac{j^2-j}{2} + 1$ and $B(j) = \frac{j^2+j}{2}$. Dividing out $\frac{1}{m}$, using $j^2 - j \geq (j-1)^2$ for $j \geq 2$, and applying the Basel problem equality, we obtain

$$\sum_{\substack{p \text{ prime}, j > 1 \\ p^j \equiv 1 \pmod{m} \\ p < m}} \frac{1}{jp^j} \leq \sum_{j=2}^{\infty} \frac{1}{(\frac{j^2-j}{2} + 1)m + 1} \leq \frac{2}{m} \cdot \sum_{j=2}^{\infty} \frac{1}{(j-1)^2} \leq \frac{\pi^2}{3m}. \quad (\text{A.116})$$

Combining Equation (A.115) and Equation (A.116), and simplifying $\pi^2/3 + 1 \leq 5$ we obtain the claim. \square

A.2.5. Estimating the Tail Part of the Main Term

Defining $\mathcal{M}_a(k) = \mathcal{M}(k)$ if $k \equiv a \pmod{m}$ and zero otherwise, and putting $\psi_a(x) = \sum_{k < x} \mathcal{M}_a(k)$, we have the following explicit result, due to Dusart [Dus98, Thm. 3.7, p. 114].

Theorem A.15 (ERH). *For every $x > \max(e^{5/4 \cdot m}, 10^{10})$, we have, assuming the Riemann Hypothesis for $L(\chi, s)$ for all Dirichlet characters χ modulo m ,*

$$|\psi_a(x) - x/\phi(m)| \leq \frac{1}{4\pi} \sqrt{x} \log^2(x)$$

Lemma A.16 (ERH). *Let m be a fixed modulus and let a be coprime with m and let $x \geq w \geq e^{5/4 \cdot m}$. Then there is a value $K_{x,w}$ that does not depend on a , and a value η_a with $|\eta_a| \leq 1$, such that*

$$\left| (S_{a,x} - S_{a,w}) - K_{x,w} - \frac{2\eta_a}{m} \right| = O(1/\log x).$$

Proof. We have

$$S_{a,x} - S_{a,w} = \sum_{\substack{p \text{ prime, } j > 0, \\ p^j \equiv a \pmod{m} \\ w < p^j \leq x}} \frac{1}{jp^j}.$$

Applying Abel summation, using that the derivative of $\frac{1}{t \log t}$ equals $\frac{-(\log(t)+1)}{\log(t)^2 t^2}$, we obtain

$$S_{a,x} - S_{a,w} = \sum_{w < k \leq x} \frac{\mathcal{M}_a(k)}{k \log k} = \frac{\psi_a(x)}{x \log x} - \frac{\psi_a(w)}{w \log w} + \int_w^x \frac{\psi_a(t)(\log(t) + 1)}{\log(t)^2 t^2} dt.$$

Writing $\psi_a(t) = \frac{t}{\phi(m)} + 1/(4\pi) \cdot \eta(t) \sqrt{t} \log^2(t)$ with $|\eta(t)| \leq 1$, we obtain that, for some η' with $|\eta'| \leq 1$,

$$\begin{aligned} & \int_w^x \frac{\psi_a(t)(\log(t) + 1)}{\log(t)^2 t^2} dt \\ &= O(1/\log(x)) + \log \log x + \log \log w - 1/\log w + \eta' \underbrace{\frac{2 \log(w) + 3}{4\pi \sqrt{w}}}_{\leq 1/m}. \end{aligned}$$

Appendix A. Appendix

Since $w \geq e^{5/4 \cdot m}$, we have $\frac{2\log(w)+3}{4\pi\sqrt{w}} \leq \frac{1}{m}$. Also, for some η'' with $|\eta''| \leq 1$, we have

$$\frac{\psi_a(w)}{w \log w} = \frac{1}{\log(w)\phi(m)} + \frac{\eta(t) \log^2(w)}{4\pi w^{1/2}} = \frac{1}{\log(w)\phi(m)} + \eta''/m$$

Combining all equations and putting $K_{x,w} = \log \log x + \log \log w - 1/\log w + \frac{1}{\log(w)\phi(m)}$, we obtain

$$\left| \sum_{w < k \leq x} \frac{\mathcal{M}_a(k)}{k \log k} - K_{x,w} - (\eta' + \eta'')/m \right| = O(1/\log(x)).$$

□

Proposition A.17 (ERH). *Let $x \geq w \geq e^{5/4 \cdot m}$. Then*

$$M_K^{(x)} - M_K^{(w)} \leq O(m/\log(x)) + 4,$$

where the implied constant is absolute (and does not depend on m).

Proof. We have, using Lemma A.16,

$$\begin{aligned} M_K^{(x)} - M_K^{(w)} &= \phi(m)(S_{1,x} - S_{1,w}) - \sum_{a \in (\mathbb{Z}/m\mathbb{Z})^*} (S_{a,x} - S_{a,w}) \\ &= (\phi(m) - \phi(m))(O(1/\log x) + K_{x,w}) + \phi(m) \cdot 2\eta_1/m + \sum_{a \in (\mathbb{Z}/m\mathbb{Z})^*} 2\eta_a/m. \end{aligned}$$

By using the fact that $|\eta_a| \leq 1$ for all $a \in (\mathbb{Z}/m\mathbb{Z})^*$, we obtain the result. □

A.3. Exact Sequences

Lemma A.18 (Kernel-cokernel exact sequence). *Let A, B, C be abelian groups and let $f : A \rightarrow B$ and $g : B \rightarrow C$ be group homomorphisms, fitting in the following commutative diagram.*

$$\begin{array}{ccc}
& B & \\
f \swarrow & \nearrow g & \\
A & \xrightarrow{gf} & C
\end{array}$$

Then, denoting ‘ker’ for the kernel of a map and ‘coker’ for the cokernel of a map, we have the following exact sequence.

$$0 \rightarrow \ker f \rightarrow \ker gf \rightarrow \ker g \rightarrow \text{coker } f \rightarrow \text{coker } gf \rightarrow \text{coker } g \rightarrow 0.$$

This sequence can be obtained mnemonically by observing the outer, blue arrows in Figure A.1.

Proof. Apply the snake lemma twice to obtain the result.

$$\begin{array}{ccccccc}
0 & \xrightarrow{\quad} & \ker f & \xrightarrow{\quad} & \ker gf & \xrightarrow{\quad} & 0 \\
\downarrow & & \downarrow & & \downarrow & & \downarrow \delta \\
0 & \xrightarrow{\quad} & A & \xrightarrow{id} & A & \xrightarrow{\quad} & 0 \\
\downarrow & & \downarrow f & & \downarrow gf & & \\
0 & \xrightarrow{\quad} & \ker g & \xrightarrow{\quad} & B & \xrightarrow{g} & C \\
\downarrow & & \downarrow & & \downarrow & & \downarrow \\
& & \ker g & \xrightarrow{\quad} & \text{coker } f & \xrightarrow{\quad} & \text{coker } gf
\end{array}$$

$$\begin{array}{ccccccc}
& \ker gf & \xrightarrow{\quad} & \ker g & \xrightarrow{\quad} & \text{coker } f & \xrightarrow{\quad} 0 \\
\downarrow & & & \downarrow & & \downarrow & \downarrow \delta \\
A & \xrightarrow{f} & B & \xrightarrow{g} & \text{coker } f & \xrightarrow{\quad} & 0 \\
\downarrow gf & & \downarrow id & & \downarrow & & \downarrow \\
C & \xrightarrow{gf} & C & \xrightarrow{\quad} & 0 & & \\
\downarrow & & \downarrow & & \downarrow & & \\
& \text{coker } gf & \xrightarrow{\quad} & \text{coker } g & \xrightarrow{\quad} & 0
\end{array}$$

□

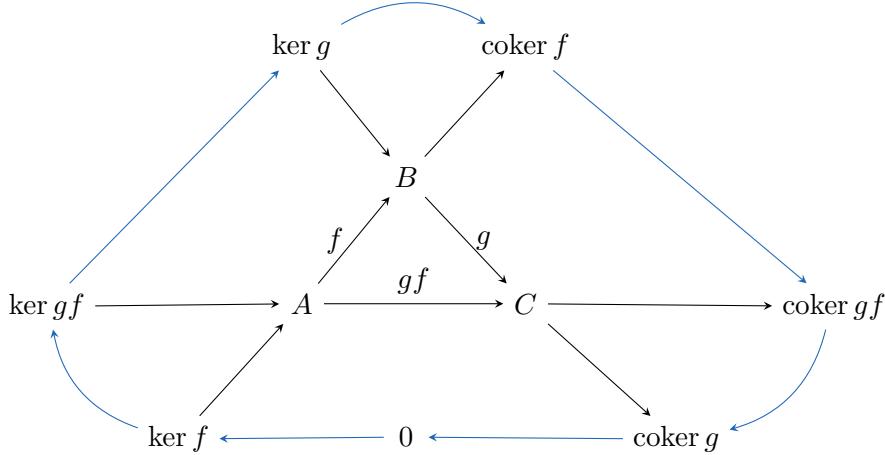


Figure A.1.: The kernel-cokernel exact sequence in the outer, blue arrows.

A.4. The Yudin-Jackson Theorem

In the chapter about the Continuous Hidden Subgroup Problem (Chapter 3), the main issue is the impact of discretization on the success probability of the quantum algorithm. This impact turns out to be largely influenced by how well a complex vector-valued function on the torus $\mathbb{T}^m = \mathbb{R}^m / \mathbb{Z}^m$ can be approximated by trigonometric functions with bounded frequencies.

This problem of finding the best trigonometric approximation has already been solved in the specific case of scalar complex functions on the torus by Yudin [Yud76], using Fourier analysis. We show here that Yudin's reasoning applies straightforwardly to vector-valued functions as well. To be clear, the following text contains the same proof as in Yudin's work [Yud76] and it is restated here for the sake of self-containedness.

Generalized result of Yudin

Recall that the L_p -norm for $p \in [1, \infty]$ for a vector-valued function $\mathbf{f} : \mathbb{T}^m \rightarrow \mathbb{C}^N$ is defined as follows¹.

$$\|\mathbf{f}\|_{p, \mathbb{T}^m} := \left(\int_{x \in \mathbb{T}^m} \|\mathbf{f}(x)\|_{\mathbb{C}^N}^p dx \right)^{1/p},$$

where $\|\cdot\|_{\mathbb{C}^N}$ is the Euclidean norm on \mathbb{C}^N . Any function for which the value $\|\mathbf{f}\|_{p, \mathbb{T}^m}$ is well-defined is called an L_p -function. For a function $\mathbf{f} : \mathbb{T}^m \rightarrow \mathbb{C}^N$ we define its Lipschitz constant to be

$$\text{Lip}(\mathbf{f}) = \inf \{L \mid \|\mathbf{f}(x) - \mathbf{f}(y)\|_{\mathbb{C}^N} \leq L \|x - y\|_{\mathbb{T}^m} \text{ for all } x, y \in \mathbb{T}^m\}.$$

For \mathbf{f} we also define a related constant, the modulus of smoothness [Yud76]:

$$\omega_2(\mathbf{f}, \delta)_p := \sup_{|y| \leq \delta} \|\mathbf{f}(\cdot - y) - 2\mathbf{f}(\cdot) + \mathbf{f}(\cdot + y)\|_{p, \mathbb{T}^m}.$$

It is evident that $\omega_2(\mathbf{f}, \delta)_p \leq \omega_2(\mathbf{f}, \delta)_\infty \leq 2 \text{Lip}(\mathbf{f}) \delta$ for functions \mathbf{f} for which both quantities are defined.

Theorem A.19 (Yudin-Jackson). *Let $\mathbf{f} : \mathbb{T}^m \rightarrow \mathbb{C}^N$ be an L_p -function. Then there exists a function $\mathbf{t} : \mathbb{T}^m \rightarrow \mathbb{C}^N$ with $\mathcal{F}_{\mathbb{T}^m}\{\mathbf{t}\}$ having support in $[-r/2, r/2]^m$ such that*

$$\|\mathbf{f} - \mathbf{t}\|_{p, \mathbb{T}^m} \leq 2\omega_2(\mathbf{f}, \sqrt{m}/r)_p \leq 2\sqrt{m} \text{Lip}(\mathbf{f})/r.$$

In essence, above theorem just states that the best trigonometric approximation of a function mainly depends on the smoothness of that function (in terms of the Lipschitz constant, for example) and how high the frequencies of the trigonometric functions are allowed to be, which is measured by r .

¹For $p = \infty$, we let $\|\mathbf{f}\|_{\infty, \mathbb{T}^m}$ to be the essential supremum of the function $x \mapsto \|\mathbf{f}\|_{\mathbb{C}^N}$.

Proof

First we prove a basic result about the modulus of smoothness; it satisfies the following ‘scaling’ property.

Lemma A.20 (Scaling property of the modulus of smoothness). *For any L_p function $\mathbf{f} : \mathbb{T}^m \rightarrow \mathbb{C}^N$ and for any $\rho, \delta > 0$, we have $\omega_2(\mathbf{f}, \rho\delta)_p \leq 2(1 + \rho^2)\omega_2(\mathbf{f}, \delta)_p$.*

Proof. Note that we have the following ‘telescopic’ finite sum

$$\begin{aligned} & \mathbf{f}(x - nt) - 2\mathbf{f}(x) + \mathbf{f}(x + nt) \\ &= \sum_{j=-n+1}^{n-1} (n - |j|)[\mathbf{f}(x + (j-1)t) - 2\mathbf{f}(x + jt) + \mathbf{f}(x + (j+1)t)]. \end{aligned}$$

So, for $|t| \leq \delta$, we have, by the triangle inequality,

$$\begin{aligned} \|\mathbf{f}(\cdot - nt) - 2\mathbf{f}(\cdot) + \mathbf{f}(\cdot + nt)\|_{p, \mathbb{T}^m} &\leq \sum_{j=-n+1}^{n-1} (n - |j|)\omega_2(\mathbf{f}, \delta)_p \\ &= n^2\omega_2(\mathbf{f}, \delta)_p. \end{aligned}$$

Therefore, for any $\rho > 0$, $\omega(\mathbf{f}, \rho\delta)_p \leq \omega(\mathbf{f}, \lceil \rho \rceil \delta)_p \leq \lceil \rho \rceil^2 \omega(\mathbf{f}, \delta)_p \leq (1 + \rho^2)\omega(\mathbf{f}, \delta)_p$. Using the fact that $(1 + \rho)^2 \leq 2(1 + \rho^2)$, we obtain the result. \square

Next, we try to approximate the function \mathbf{f} by the function $\mathbf{f} \star K$, a convolution of f with a suitable kernel K . The closeness of this approximation largely depends on the smoothness of \mathbf{f} and the value of a certain integral involving the kernel K .

Lemma A.21. *Let $K : \mathbb{T}^m \rightarrow [0, \infty)$ be a L_1 -function satisfying $\int_{t \in \mathbb{T}^m} K(t)dt = 1$ and $K(-t) = t$ for all $t \in \mathbb{T}^m$. Denote $\mathbf{t} = \mathbf{f} \star K = \int_{t \in \mathbb{T}^m} \mathbf{f}(\cdot - t)K(t)dt$. Then, for all $r > 0$,*

$$\|\mathbf{f} - \mathbf{t}\|_{p, \mathbb{T}^m} \leq \omega_2(\mathbf{f}, \sqrt{m}/r)_p \left(1 + \frac{r^2}{m} \int_{t \in [-1/2, 1/2]^m} |t|^2 \cdot K(t)dt \right), \quad (\text{A.117})$$

Proof. By the fact that K is even,

$$\begin{aligned}\mathbf{t}(x) = \mathbf{f} \star K(x) &= \int_{t \in \mathbb{T}^m} \mathbf{f}(x-t)K(t)dt = \int_{t \in \mathbb{T}^m} \mathbf{f}(x+t)K(t)dt \\ &= \frac{1}{2} \int_{t \in \mathbb{T}^m} \mathbf{f}(x-t) + \mathbf{f}(x+t)K(t)dt.\end{aligned}$$

We can write $f(x) = \int_{t \in \mathbb{T}^m} f(x)K(t)dt$, since $\int_{t \in \mathbb{T}^m} K(t)dt = 1$. Therefore,

$$\mathbf{t}(x) - \mathbf{f}(x) = \frac{1}{2} \int_{t \in \mathbb{T}^m} (\mathbf{f}(x-t) - 2\mathbf{f}(x) + \mathbf{f}(x+t))K(t)dt.$$

Taking L_p -norms, using the integral-triangle inequality, integrating over the set $[-1/2, 1/2]^m$, using the fact that $K(t)$ is a positive scalar and applying Lemma A.20 with $\delta = \sqrt{m}/r$ and $\rho = r|t|/\sqrt{m}$, we obtain

$$\begin{aligned}\|\mathbf{f} - \mathbf{t}\|_{p, \mathbb{T}^m} &\leq \frac{1}{2} \int_{t \in [-1/2, 1/2]^m} \omega_2(\mathbf{f}, |t|)_p K(t)dt \\ &\leq \int_{t \in [-1/2, 1/2]^m} \left(1 + \frac{|t|^2 r^2}{m}\right) \omega_2(\mathbf{f}, \sqrt{m}/r)_p K(t)dt.\end{aligned}$$

Rewriting the integral, using $\int_{t \in \mathbb{T}^m} K(t)dt = 1$, we arrive at Equation (A.117). \square

In the next step, we will instantiate the kernel $K = K_r$ in such a way that its Fourier coefficients have support in $[-r/2, r/2]^m$. This means, by the convolution formula, that $\mathbf{t} = \mathbf{f} \star K_r$ also has Fourier coefficients with support only in $[-r/2, r/2]^m$. Furthermore, K_r is chosen in such a way that

$$\frac{r^2}{m} \cdot \int_{t \in [-1/2, 1/2]^m} |t|^2 K_r(t)dt \leq 1.$$

Lemma A.22. *Let $\lambda = \phi \star \phi = \int_{t \in \mathbb{R}^m} \phi(\cdot - t)\phi(t)dt$, where*

$$\phi(x_1, \dots, x_m) = \begin{cases} 2^m \prod_{j=1}^m \cos(2\pi x_j) & \text{if } (x_1, \dots, x_m) \in [-1/4, 1/4]^m \\ 0 & \text{otherwise} \end{cases}$$

Furthermore, define $K_r : \mathbb{T}^m \rightarrow \mathbb{C}$ by the rule $K_r(t) := \mathcal{F}_{\mathbb{T}^m}^{-1}\{\lambda(\cdot/r)|_{\mathbb{Z}^m}\}(t) = \sum_{z \in \mathbb{Z}^m} \lambda(z/r) e^{2\pi i \langle t, z \rangle}$. Then

- (i) $K_r(t) \geq 0$ and $K_r(t) = K_r(-t)$ for all $t \in \mathbb{T}^m$,
- (ii) $\int_{t \in \mathbb{T}^m} K_r(t) dt = 1$,
- (iii) $\mathcal{F}_{\mathbb{T}^m}\{K_r\}$ has support only in $[-r/2, r/2]^m$,
- (iv) $\int_{t \in \mathbb{T}^m} |t|^2 K_r(t) dt \leq m/r^2$.

Proof. For (i), note that K_r is even because λ is. For positivity, we apply the Poisson summation formula.

$$K_r = \mathcal{F}_{\mathbb{T}^m}^{-1}\{\lambda(\cdot/r)|_{\mathbb{Z}^m}\} = \mathcal{F}_{\mathbb{R}^m}^{-1}\{\lambda(\cdot/r)\}|^{\mathbb{Z}^m} = r^m \hat{\lambda}(r\cdot)|^{\mathbb{Z}^m} \geq 0.$$

The last inequality follows from the convolution formula: $\hat{\lambda} = \widehat{\phi \star \phi} = \hat{\phi} \cdot \hat{\phi} \geq 0$. For (ii), note that $\int_{t \in \mathbb{T}^m} K_r(t) dt = \mathcal{F}_{\mathbb{T}^m}\{K_r\}[0] = \lambda(0) = \int_{t \in \mathbb{R}^m} \phi(t)^2 dt = 1$. Part (iii) is can be shown by combining the following facts: $\mathcal{F}_{\mathbb{T}^m}\{K_r\} = \lambda(\cdot/r)|_{\mathbb{Z}^m}$ and $\lambda(x) = 0$ if $|x|_\infty > 1/2$. Part (iv) is the most technical; since $K_r = r^m \hat{\lambda}(r\cdot)|^{\mathbb{Z}^m}$ and $|t|^2 \leq |t+v|^2$ for any $v \in \mathbb{Z}^m$ and $t \in [-1/2, 1/2]^m$, we have

$$\begin{aligned} \int_{t \in [-\frac{1}{2}, \frac{1}{2}]^m} |t|^2 K_r(t) dt &= \int_{t \in [-\frac{1}{2}, \frac{1}{2}]^m} |t|^2 r^m \sum_{z \in \mathbb{Z}^m} \hat{\lambda}(r(t+z)) dt \\ &\leq \int_{\mathbb{R}^m} |t|^2 \hat{\lambda}(rt) r^m dt = r^{-2} \int_{\mathbb{R}^m} |y|^2 \hat{\lambda}(y) dy, \end{aligned} \quad (\text{A.118})$$

where the last equality holds by the substitution rule. By the definition of λ , Plancherel's theorem and the fact that $2\pi i y \hat{\phi} = \mathcal{F}_{\mathbb{R}^m}\{\nabla \phi\}$, we obtain that the right side of Equation (A.118) equals

$$r^{-2} \int_{y \in \mathbb{R}^m} |y|^2 \hat{\phi}(y) \hat{\phi}(y) dy = r^{-2} \|y \hat{\phi}(y)\|_{2, \mathbb{R}^m}^2 = r^{-2} \|(2\pi)^{-1} \nabla \phi\|_{2, \mathbb{R}^m}^2 = m/r^2.$$

where the last equation follows from integrating the following function over \mathbb{R}^m , which proves (iv).

$$|(2\pi)^{-1} \nabla \phi(x)|^2 = \begin{cases} 2^{2m} \sum_{j=1}^m \sin^2(2\pi x_j) \prod_{k \neq j} \cos^2(2\pi x_k) & \text{if } \mathbf{x} \in [-\frac{1}{4}, \frac{1}{4}]^m \\ 0 & \text{otherwise} \end{cases}$$

□

Combining Lemma A.22 and Lemma A.21 we arrive at a proof for Theorem A.19.

Proof of Theorem A.19. Put $\mathbf{t} = \mathbf{f} \star K_r = \int_{t \in \mathbb{T}^m} \mathbf{f}(t) K_r(\cdot - t) dt$ with K_r as in Lemma A.22. As K_r satisfies the requirements of Lemma A.21 and

$$\frac{r^2}{m} \int_{t \in [-1/2, 1/2]^m} |t|^2 K_r(t) dt \leq 1,$$

by Lemma A.22(iv), we have $\|\mathbf{f} - \mathbf{t}\|_{p, \mathbb{T}^m} \leq 2\omega_2(\mathbf{f}, \sqrt{m}/r) \leq 2\sqrt{m} \text{Lip}(f)/r$. By Lemma A.22(iii) and the convolution formula, we have $\mathcal{F}_{\mathbb{T}^m}\{\mathbf{t}\} = \mathcal{F}_{\mathbb{T}^m}\{\mathbf{f}\} \cdot \mathcal{F}_{\mathbb{T}^m}\{K_r\} = \mathcal{F}_{\mathbb{T}^m}\{\mathbf{f}\} \cdot \lambda(\cdot/r)|_{\mathbb{Z}^m}$. Since $\lambda(\cdot/r)$ only has support in $[-r/2, r/2]^m$, the Fourier transform of \mathbf{t} has also only support there. \square

A.5. The Gaussian State

A.5.1. Reducing to the One-dimensional Case

In this section, we estimate the exact quantum complexity of obtaining an approximation, in the trace distance, of the state

$$\frac{1}{\sqrt{\rho_{1/s}(\mathbb{D}_{\text{rep}}^m)}} \sum_{\mathbf{x} \in \mathbb{D}_{\text{rep}}^m} \sqrt{\rho_{1/s}(\mathbf{x})} |\mathbf{x}\rangle, \quad (\text{A.119})$$

where $\mathbb{D}_{\text{rep}}^m = \frac{1}{q} \mathbb{Z}^m \cap [-1/2, 1/2]^m$, and where $\rho_{1/s}(\cdot) = e^{-\pi s^2 \|\cdot\|^2}$ is the Gaussian function (see Section 2.5.3).

An element $|\mathbf{x}\rangle$ with $\mathbf{x} = (x_1, \dots, x_m) \in \mathbb{D}_{\text{rep}}^m$ is represented as a tensor product $|x_1\rangle \otimes \dots \otimes |x_m\rangle$. As the function $\sqrt{\rho_{1/s}(x)} = \rho_{\sqrt{2}/s}(x)$ can be written as a product of functions with separated variables as well, we obtain that Equation (A.119) equals

$$\bigotimes_{j=1}^m \frac{1}{\sqrt{\rho_{1/s}(\frac{1}{q}[q]_c)}} \sum_{x \in \frac{1}{q}[q]_c} \sqrt{\rho_{1/s}(x)} |x\rangle,$$

where $\frac{1}{q}[q]_c = \frac{1}{q}\mathbb{Z} \cap [-1/2, 1/2]$. Therefore, the problem of approximating the state as in Equation (A.119) reduces to the one-dimensional case. By

rescaling the variable $x \in \frac{1}{q}[q]_c$, the computation of this one-dimensional state boils down to calculating the following quantum state, with $\varsigma = q/s$.

$$|\rho_{\varsigma,q}\rangle := \frac{1}{\sqrt{\rho_{\varsigma}([q]_c)}} \sum_{x \in [q]_c} \sqrt{\rho_{\varsigma}(x)} \cdot |x\rangle.$$

Here, $[q]_c = \{-\frac{q}{2} + 1, \dots, 0, \dots, \frac{q}{2}\}$, and $q = 2^Q$ is a 2-power, for simplicity.

A.5.2. The Periodic and Non-periodic Discrete Gaussian

To obtain a Gaussian superposition in one dimension, we follow a method of Kitaev and Webb [KW08]. Their algorithm is an improvement of that of Grover and Rudolph [GR02].

Kitaev and Webb's algorithm actually does not compute a discrete Gaussian quantum state, but something very close; a *periodized* discrete Gaussian quantum state. This periodized state has the advantage of having a more natural normalization and, more importantly, having a specific *sum decomposition*. These advantages lead to a slightly more efficient algorithm [KW08] computing the discrete Gaussian superposition, compared to the algorithm of Grover and Rudolph.

Definition A.23 (Discrete Periodized Gaussian function). *For $\varsigma \in \mathbb{R}_{>0}$ and $q = 2^Q$ a power of two, we denote by $\xi_{\varsigma,q} : \mathbb{Z}/q\mathbb{Z} \rightarrow \mathbb{R}_{>0}$ the function defined by the following rule*

$$\xi_{\varsigma,q}(x) = \sqrt{\sum_{z \in \mathbb{Z}} \rho_{\varsigma}(x + qz)}.$$

The associated quantum state is defined as follows

$$|\xi_{\varsigma,q}\rangle = \frac{1}{\sqrt{\rho_{\varsigma}(\mathbb{Z})}} \sum_{x \in [q]_c} \xi_{\varsigma,q}(x) |x\rangle$$

Lemma A.24. *Let $\varsigma \in \mathbb{R}_{>0}$ and $q = 2^Q \in \mathbb{N}$, with $q \geq \varsigma$. Then*

$$D(|\xi_{\varsigma,q}\rangle, |\rho_{\varsigma,q}\rangle) \leq \exp\left(-\frac{q^2}{2\varsigma^2}\right)$$

where D is the trace distance [NC11, §9.2.1].

Proof. Since $\xi_{\varsigma,q}(x) \geq \sqrt{\rho_{\varsigma}(x)}$, we have, writing out the definitions,

$$\begin{aligned} \langle \xi_{\varsigma,q} | \rho_{\varsigma,q} \rangle &= \frac{\sum_{x \in [q]_c} \xi_{\varsigma,q}(x) \sqrt{\rho_{\varsigma}(x)}}{\sqrt{\rho_{\varsigma}(\mathbb{Z}) \rho_{\varsigma}([q]_c)}} \\ &\geq \frac{\sum_{x \in [q]_c} \rho_{\varsigma}(x)}{\sqrt{\rho_{\varsigma}(\mathbb{Z}) \rho_{\varsigma}([q]_c)}} = \sqrt{\rho_{\varsigma}([q]_c) / \rho_{\varsigma}(\mathbb{Z})}. \end{aligned}$$

Since the trace distance between the pure states $|\xi_{\varsigma,q}\rangle$ and $|\rho_{\varsigma,q}\rangle$ is equal to $\sqrt{1 - |\langle \xi_{\varsigma,q} | \rho_{\varsigma,q} \rangle|^2}$ [NC11, §9.2], we obtain

$$\begin{aligned} D(|\xi_{\varsigma,q}\rangle, |\rho_{\varsigma,q}\rangle) &\leq \sqrt{1 - \rho_{\varsigma}([q]_c) / \rho_{\varsigma}(\mathbb{Z})} = \sqrt{\rho_{\varsigma}(\mathbb{Z} \setminus [q]_c)} \\ &\leq \sqrt{\beta_{q/\varsigma}^{(1)}} \leq \exp\left(-\frac{q^2}{2\varsigma^2}\right), \end{aligned}$$

where we applied Banaszczyk's tail bound (see Lemma 2.25). \square

Above lemma essentially states that whenever q is relatively large, and ς is not too large, then the periodic discrete Gaussian and the (non-periodic) discrete Gaussian are very close in trace distance. That has as a consequence that the associated measurement probability distributions are close in total variation distance [NC11, Thm. 9.1].

A.5.3. Computing the Periodic Gaussian State

According to the previous subsection, we can resort to computing the state $|\xi_{\varsigma,q}\rangle$ instead of $|\rho_{\varsigma,q}\rangle$, as they are close to each other for a suitable choice of parameters. As already mentioned, the quantum state $|\xi_{\varsigma,q}\rangle$ can be decomposed into a superposition that can be exploited algorithmically. In order to phrase this decomposition we first introduce the following notation of a quantum state 'translated' by $t \in \mathbb{R}$.

$$|\xi_{\varsigma,q}(\cdot + t)\rangle = \frac{1}{\sqrt{\rho_{\varsigma}(\mathbb{Z} + t)}} \sum_{x \in [q]_c} \xi_{\varsigma,q}(x + t) |x\rangle$$

Likewise, we denote

$$|\rho_{\varsigma,q}(\cdot + t)\rangle := \frac{1}{\sqrt{\rho_{\varsigma}([q]_c + t)}} \sum_{x \in [q]_c} \sqrt{\rho_{\varsigma}(x + t)} \cdot |x\rangle$$

Now we are ready to state the decomposition lemma.

Lemma A.25 ([KW08, Eq. (11)]). *Let $\varsigma \in \mathbb{R}_{>0}$, $t \in \mathbb{R}$ and let $q \in \mathbb{N}$ be even. Then*

$$|\xi_{\varsigma,q}(\cdot + 2t)\rangle = |\xi_{\frac{\varsigma}{2}, \frac{q}{2}}(\cdot + t)\rangle \otimes \cos \alpha |0\rangle + |\xi_{\frac{\varsigma}{2}, \frac{q}{2}}(\cdot + t + \frac{1}{2})\rangle \otimes \sin \alpha |1\rangle,$$

with $\alpha = \arccos \left(\sqrt{\rho_{\frac{\varsigma}{2}}(\mathbb{Z} + t) / \rho_{\varsigma}(\mathbb{Z} + 2t)} \right)$.

Proof. Splitting the sum into a part with even numbers and a part with odd numbers, we obtain

$$\begin{aligned} & \sqrt{\rho_{\varsigma}(\mathbb{Z} + 2t)} \cdot |\xi_{\varsigma,q}(\cdot + 2t)\rangle \\ &= \sum_{x \in [q]_c} \xi_{\varsigma,q}(x + 2t) |j\rangle \\ &= \sum_{x \in [\frac{q}{2}]_c} \xi_{\varsigma,q}(2x + 2t) |x\rangle |0\rangle + \sum_{x \in [\frac{q}{2}]_c} \xi_{\varsigma,q}(2x + 1 + 2t) |x\rangle |1\rangle. \end{aligned} \quad (\text{A.120})$$

We now focus the computation on the sum over the odd numbers, as the computation for the even numbers is similar. By writing out the definition of $\xi_{\varsigma,q}(x)$ and putting the scalar 2 into the standard deviation ς , we obtain

$$\begin{aligned} \xi_{\varsigma,q}(2x + 1 + 2t)^2 &= \rho_{\varsigma}(2x + 1 + 2t + q\mathbb{Z}) \\ &= \rho_{\frac{\varsigma}{2}}(x + t + \frac{1}{2} + \frac{q}{2} \cdot \mathbb{Z}) = \xi_{\frac{\varsigma}{2}, \frac{q}{2}}(x + \frac{1}{2} + t)^2. \end{aligned}$$

Using a similar computation for the even case and writing out the definitions, we obtain

$$\begin{aligned} & \sqrt{\rho_{\varsigma}(\mathbb{Z} + 2t)} \cdot |\xi_{\varsigma,q}(\cdot + 2t)\rangle \\ &= \sqrt{\rho_{\frac{\varsigma}{2}}(\mathbb{Z} + t)} \cdot |\xi_{\frac{\varsigma}{2}, \frac{q}{2}}(\cdot + t)\rangle \otimes |0\rangle + \sqrt{\rho_{\frac{\varsigma}{2}}(\mathbb{Z} + t + \frac{1}{2})} \cdot |\xi_{\frac{\varsigma}{2}, \frac{q}{2}}(\cdot + t + \frac{1}{2})\rangle \otimes |1\rangle. \end{aligned}$$

Dividing above expression by $\sqrt{\rho_{\varsigma}(\mathbb{Z} + 2t)}$ we obtain Equation (A.120), where we use the fact that $\rho_{\varsigma/2}(\mathbb{Z} + t) + \rho_{\varsigma/2}(\mathbb{Z} + t + \frac{1}{2}) = \rho_{\varsigma}(\mathbb{Z} + 2t)$. \square

This lemma directly leads to an algorithm for computing (an approximation of) the state $|\xi_{\varsigma,q}\rangle$, which is spelled out in Algorithm 10.

Algorithm 10: Recursive algorithm preparing the periodic Gaussian state

Require: The parameters $\varsigma \in \mathbb{R}_{>0}$, $t \in \mathbb{R}$, $k \in \mathbb{N}$ and $q = 2^Q \in \mathbb{N}$.

Ensure: An approximation of the state $|\xi_{\varsigma,q}(\cdot + t)\rangle$

- 1: **Initial state:** $|t, \varsigma, q\rangle|0^Q\rangle$;
- 2: **Compute the α -rotation by on the last qubit:** Compute α with bit-precision k and store it in a k -qubit ancilla register. Apply the α -rotation on the last qubit and uncompute α again, which yields the state $|t, \varsigma, q\rangle|0^{Q-1}\rangle(\cos \alpha|0\rangle + \sin \alpha|1\rangle)$;
- 3: **Apply a parameter change, controlled by the last qubit** yielding $\cos \alpha|\frac{t}{2}, \frac{\varsigma}{2}, \frac{q}{2}\rangle|0^{Q-1}\rangle|0\rangle + \sin \alpha|\frac{t+1}{2}, \frac{\varsigma}{2}, \frac{q}{2}\rangle|0^{Q-1}\rangle|1\rangle$;
- 4: **Apply quantum recursion (step 2 and 3) on all qubits except the last, whenever $q > 1$, yielding**
 $\cos \alpha|\frac{t}{2}, \frac{\varsigma}{2}, \frac{q}{2}\rangle|\xi_{\frac{\varsigma}{2}, \frac{q}{2}}(\cdot + \frac{t}{2})\rangle|0\rangle + \sin \alpha|\frac{t+1}{2}, \frac{\varsigma}{2}, \frac{q}{2}\rangle|\xi_{\frac{\varsigma}{2}, \frac{q}{2}}(\cdot + \frac{t+1}{2})\rangle|0\rangle$;
- 5: **Un-apply the controlled parameter change, yielding**
 $|t, \varsigma, q\rangle \left(\cos \alpha|\xi_{\frac{\varsigma}{2}, \frac{q}{2}}(\cdot + \frac{t}{2})\rangle|0\rangle + \sin \alpha|\xi_{\frac{\varsigma}{2}, \frac{q}{2}}(\cdot + \frac{t+1}{2})\rangle|1\rangle \right) =$
 $|t, \varsigma, q\rangle|\xi_{\varsigma,q}(\cdot + t)\rangle$;

A.5.4. Estimating the Complexity and Fidelity of Algorithm 10

We will discuss now how well Algorithm 10 approximates the state $|\xi_{\varsigma,q}\rangle$. For ease of analysis, we will assume (without loss of generality) that the operations on the parameters ς (in step 3 of Algorithm 10) are exact. Then it turns out that the approximation error is primarily caused by the fact that the angle α in the algorithm is computed up to bit precision k (meaning, with error at most 2^{-k}). This is made precise in the following lemma.

Lemma A.26. *Let $|\tilde{\xi}_{\varsigma,q}(\cdot + t)\rangle$ be the output of Algorithm 10 with input parameters $\varsigma \in \mathbb{R}_{>0}$, $k \in \mathbb{N}$, $q = 2^Q \in \mathbb{N}$ and $t \in (-1, 1)$, then we have*

$$T\left(|\tilde{\xi}_{\varsigma,q}(\cdot + t)\rangle, |\xi_{\varsigma,q}(\cdot + t)\rangle\right) \leq 2^{-k}Q$$

where T denotes the trace distance.

Proof. The proof proceeds by induction on Q , where $q = 2^Q$. We use the the identity $D(|\psi\rangle, |\phi\rangle)^2 + |\langle\psi|\phi\rangle|^2 = 1$ multiple times throughout the proof (see [NC11, §9.2]). Let $\tilde{\alpha}$ be a k -bit approximation of α , i.e., $|\alpha - \tilde{\alpha}| < 2^{-k}$, and denote $|\tilde{\xi}_{\varsigma,q}(\cdot + t)\rangle = \cos \tilde{\alpha} |\tilde{\xi}_{\frac{\varsigma}{2}, \frac{q}{2}}(\cdot + \frac{t}{2})\rangle |0\rangle + \sin \tilde{\alpha} |\tilde{\xi}_{\frac{\varsigma}{2}, \frac{q}{2}}(\cdot + \frac{t+1}{2})\rangle |1\rangle$ for the output of Algorithm 10 with input parameters $\varsigma, k, q = 2^Q$ and $t \in (-1, 1)$. Without loss of generality, we assume that $t = 0$ for sake of clarity; for arbitrary $t \in (-1, 1)$ the calculation is similar.

$$\langle \tilde{\xi}_{\varsigma,q} | \xi_{\varsigma,q} \rangle = \cos(\alpha) \cos(\tilde{\alpha}) \langle \tilde{\xi}_{\frac{\varsigma}{2}, \frac{q}{2}} | \xi_{\frac{\varsigma}{2}, \frac{q}{2}} \rangle + \sin(\alpha) \sin(\tilde{\alpha}) \langle \tilde{\xi}_{\frac{\varsigma}{2}, \frac{q}{2}}(\cdot + \frac{1}{2}) | \xi_{\frac{\varsigma}{2}, \frac{q}{2}}(\cdot + \frac{1}{2}) \rangle.$$

By the induction hypothesis, we have

$$|\langle \tilde{\xi}_{\frac{\varsigma}{2}, \frac{q}{2}}(\cdot + t) | \xi_{\frac{\varsigma}{2}, \frac{q}{2}}(\cdot + t) \rangle| \geq \sqrt{1 - (Q - 1)^2 2^{-2k}}$$

for $t \in (-1, 1)$. Using the trigonometric identity $\cos(\alpha) \cos(\tilde{\alpha}) + \sin(\alpha) \sin(\tilde{\alpha}) = \cos(\alpha - \tilde{\alpha})$ and the fact that the periodic Gaussian state only has positive amplitudes, we obtain

$$|\langle \tilde{\xi}_{\varsigma,q} | \xi_{\varsigma,q} \rangle| \geq \cos(\alpha - \tilde{\alpha}) \sqrt{1 - (Q - 1)^2 2^{-2k}}$$

Therefore $D(|\xi_{\varsigma,q}\rangle, |\tilde{\xi}_{\varsigma,q}\rangle) = \sqrt{1 - |\langle \xi_{\varsigma,q} | \tilde{\xi}_{\varsigma,q} \rangle|^2} \leq \sin(\alpha - \tilde{\alpha}) + (Q - 1)2^{-k} \leq Q2^{-k}$. Note that we omitted the base case, which can be done by a very similar computation using the same trigonometric identity. \square

Lemma A.27. *Computing α with k -bits of precision in step 2 of Algorithm 10 can be done within $O(k^{3/2} \cdot \text{polylog}(k))$ operations.*

Proof. Can be found in Appendix A.5.5. \square

Proposition A.28. *Algorithm 10 with input $\varsigma \in \mathbb{R}_{>0}$, $k \in \mathbb{N}$, $q = 2^Q \in \mathbb{N}$ and $t \in (-1, 1)$ uses $O(Q + k)$ qubits and $O(Q \cdot k^{3/2} \cdot \text{polylog}(k))$ quantum gates.*

Proof. The number of qubits used in Algorithm 10 equals $O(Q + k)$, because α is stored in k ancilla qubits during step 2 with bit precision k . The variable $\varsigma \in \mathbb{R}$ can be stored with similar precision.

For the number of gates, we go through the relevant steps of Algorithm 10. Step 2 computes (and uncomputes) α with precision 2^{-k} . By Lemma A.27, this costs at most $O(k^{3/2} \cdot \text{polylog}(k))$ quantum gates. The α -rotation in this step costs k quantum gates, as a sequence of controlled $R_{\pi/2^j}$ -gates.

Step 3 (and step 5) is a parameter change, which costs a mere constant number of gates. Step 6 applies recursion, which, by induction, costs $O((Q - 1) \cdot k^{3/2} \cdot \text{polylog}(k))$ gates. Adding all together gives a number of $O(Q \cdot k^{3/2} \cdot \text{polylog}(k))$ gates. \square

Theorem A.29. *For $q = 2^Q \in \mathbb{N}$, $k \in \mathbb{N}$ and $\varsigma > 1$, there exists an quantum algorithm that prepares the one-dimensional Gaussian state*

$$|\rho_{\varsigma,q}\rangle = \frac{1}{\sqrt{\rho_{\varsigma}([q]_c)}} \cdot \sum_{x \in [q]_c} \sqrt{\rho_{\varsigma}(x)} |x\rangle \quad (\text{A.121})$$

within trace distance $\exp(-\frac{q^2}{2\varsigma^2}) + \log(q)2^{-k}$, using $O(\log(q) + k)$ qubits and $O(\log(q) \cdot k^{3/2} \cdot \text{polylog}(k))$ quantum gates. Here, $[q]_c$ denotes $\{-\frac{q}{2}, \dots, \frac{q-1}{2}\}$.

Proof. The state in Equation (A.121) can be approximated by running Algorithm 10 with parameters $\varsigma, q = 2^Q$, $t = 0$ and k . Combining Lemma A.24 and Lemma A.26 and using the fact that we can add trace distances [NC11, Ch. 9], this approximation is within trace distance $\exp(-\frac{q^2}{2\varsigma^2}) + Q2^{-k}$.

For the running time, use Proposition A.28 to conclude that Algorithm 10 with the mentioned parameters uses $O(Q + k)$ qubits and $O(Q \cdot k^{3/2})$ quantum gates, which proves the claim. \square

Theorem 3.12. *For $q = 2^Q \in \mathbb{N}$, error parameter $\eta \in (0, 1)$ and $s > 2\sqrt{\log(m/\eta)}$, there exists a quantum algorithm that prepares the higher-dimensional Gaussian state*

$$\frac{1}{\sqrt{\rho_{1/s}(\mathbb{D}_{\text{rep}}^m)}} \sum_{\mathbf{x} \in \mathbb{D}_{\text{rep}}^m} \sqrt{\rho_{1/s}(\mathbf{x})} |\mathbf{x}\rangle = \bigotimes_{j=1}^m \frac{1}{\sqrt{\rho_{1/s}(\frac{1}{q}[q]_c)}} \sum_{x \in \frac{1}{q}[q]_c} \sqrt{\rho_{1/s}(x)} |x\rangle,$$

within trace distance η , using $O(mQ + \log(\eta^{-1}))$ qubits and using $O(mQ \cdot \log(mQ\eta^{-1})^2)$ quantum gates.

Proof. Instantiating Theorem A.29 with $\varsigma = q/s$ and $k = \lceil \log(2mQ\eta^{-1}) \rceil$ and rescaling the states x by q , gives the desired quantum state.

Note that the trace distance needs to be multiplied by m , due to the m -fold tensor product. This yields a trace distance of $m \exp(-s^2/2) + mQ2^{-k} \leq \frac{1}{2}\eta + \frac{1}{2}\eta \leq \eta$. Regarding qubits, we need $O(mQ)$ qubits for storing the m -dimensional Gaussian state and $O(k) = O(\log(\eta^{-1}) + \log(mQ))$ ancilla qubits, for computing and uncomputing the rotation angle α . Together this is at most $O(mQ + \log(\eta^{-1}))$ qubits.

For the number of quantum gates we just multiply the number of gates used in Theorem A.29 by m , instantiating $k = \lceil \log(2mQ\eta^{-1}) \rceil$ and simplifying the expressions using the big-O notation:

$$O(m \cdot \log(q) \cdot k^{3/2} \cdot \text{polylog}(k)) \leq O(mQ \cdot k^2) = O(mQ \cdot \log(mQ\eta^{-1})^2).$$

□

A.5.5. Proof of Lemma A.27

Lemma A.30. *The value $\rho_{\frac{\mu}{2}, \frac{\varsigma}{2\sqrt{2}}}(\mathbb{Z})$ can be computed with relative precision 2^{-k} within time $O(k^{3/2} \text{polylog}(k))$.*

Proof. We distinguish two cases.

- $\varsigma < \sqrt{2}$. Then, by Lemma 2.25,

$$\left| \rho_{\mu, \frac{\varsigma}{\sqrt{2}}}(\mathbb{Z}) - \rho_{\lfloor \mu \rfloor, \frac{\varsigma}{\sqrt{2}}}(\{-h, \dots, 0, \dots, h\}) \right| \leq \beta_{\sqrt{2}h/\varsigma}^{(1)} \cdot \rho_{\mu, \frac{\varsigma}{\sqrt{2}}}(\mathbb{Z}).$$

- $\varsigma > \sqrt{2}$. Applying the Poisson summation formula, we obtain

$$\rho_{\mu, \frac{\varsigma}{\sqrt{2}}}(\mathbb{Z}) = \frac{\varsigma}{\sqrt{2}} \sum_{t \in \mathbb{Z}} \rho_{0, \frac{\sqrt{2}}{\varsigma}}(t) e^{-2\pi i t \mu}.$$

Therefore

$$\left| \rho_{\mu, \frac{\varsigma}{\sqrt{2}}}(\mathbb{Z}) - \frac{\varsigma}{\sqrt{2}} \sum_{t \in \{-h, \dots, 0, \dots, h\}} \rho_{\frac{\sqrt{2}}{\varsigma}}(t) e^{-2\pi i t \mu} \right| \leq \frac{\varsigma}{\sqrt{2}} \beta_{\varsigma h / \sqrt{2}}^{(1)} \cdot \rho_{0, \sqrt{2}/\varsigma}(\mathbb{Z})$$

which is bounded by $\beta_{\varsigma h / \sqrt{2}}^{(1)} \cdot \rho_{0, \varsigma / \sqrt{2}}(\mathbb{Z}) \leq 2\beta_{\varsigma h / \sqrt{2}}^{(1)} \cdot \rho_{\mu, \frac{\varsigma}{\sqrt{2}}}(\mathbb{Z})$, by the Poisson summation formula and by smoothing arguments (see Lemma 2.31), as $\rho_{\mu, \varsigma / \sqrt{2}}(\mathbb{Z}) \geq (1 - 2\beta_{s/\sqrt{2}}^{(1)}) \rho_{0, \varsigma / \sqrt{2}} \geq \frac{1}{2} \rho_{0, \varsigma / \sqrt{2}}$.

So the relative error is at most $2\beta_h^{(1)} \leq e^{-(h-1)^2}$ for $h > 2$. Therefore, choosing $h = k^{1/2} + 1$ is enough to compute $\rho_{\frac{\mu}{2}, \frac{\varsigma}{2\sqrt{2}}}(\mathbb{Z})$ with relative error 2^{-k} . Because evaluating an exponential function takes $O(k \cdot \text{polylog}(k))$ time [Bre10], we arrive at the claim. \square

Lemma A.31. *The fraction $\rho_{\frac{\mu}{2}, \frac{\varsigma}{2\sqrt{2}}}(\mathbb{Z}) / \rho_{\mu, \frac{\varsigma}{\sqrt{2}}}(\mathbb{Z})$ can be computed with precision 2^{-k} within time $O(k^{3/2} \cdot \text{polylog}(k))$.*

Proof. Denote $a = \rho_{\frac{\mu}{2}, \frac{\varsigma}{2\sqrt{2}}}(\mathbb{Z})$ and $b = \rho_{\mu, \frac{\varsigma}{\sqrt{2}}}(\mathbb{Z})$. Suppose we have relative errors $|\tilde{a} - a| \leq 2^{-k}a/2 \leq 2^{-k}b/2$, $|\tilde{b} - b| \leq 2^{-k}b/2$ and $\tilde{a}/\tilde{b} < 1$, then $\left| \frac{\tilde{a}}{\tilde{b}} - \frac{a}{b} \right| \leq \frac{|\tilde{b}(a-\tilde{a}) - \tilde{a}(b-\tilde{b})|}{b\tilde{b}} \leq \frac{|a-\tilde{a}|}{b} + \frac{|b-\tilde{b}|}{b} \leq 2^{-k}$. By Lemma A.30, we see that both a and b can be computed within relative precision $2^{-k}/2$ within time $O(k^{3/2} \text{polylog}(k))$. Therefore, the fraction a/b can be computed with absolute precision 2^{-k} within time $O(k^{3/2} \text{polylog}(k))$. \square

Lemma A.32. *For $x \in [0, 1 - \varepsilon]$ and $\varepsilon < \frac{3}{4}$, we have*

$$|\arccos(\sqrt{x + \varepsilon}) - \arccos(\sqrt{x})| \leq 8\sqrt{\varepsilon}$$

Proof. The derivative of $\arccos(\sqrt{t})$ equals $w(t) = -\frac{2}{\sqrt{(1-t)t}}$. Therefore

$$\begin{aligned} |\arccos(\sqrt{x+\varepsilon}) - \arccos(\sqrt{x})| &\leq \left| \int_x^{x+\varepsilon} w(t) dt \right| \\ &\leq \int_x^{x+\varepsilon} |w(t)| dt \leq \int_0^\varepsilon |w(t)| dt. \end{aligned}$$

The last inequality follows from the fact that $w(t)$ is both strictly decreasing on $[0, 1/2]$ and symmetric around $t = 1/2$. The claim then follows from the bound $\int_0^\varepsilon |w(t)| dt = \int_0^\varepsilon \frac{2}{\sqrt{(1-x)x}} \leq 4 \int_0^\varepsilon \frac{dt}{\sqrt{t}} = 8\sqrt{\varepsilon}$. \square

By combining Lemma A.31 and Lemma A.32, we obtain that the expression $\arccos \sqrt{\rho_{\frac{\mu}{2}, \frac{\varsigma}{2\sqrt{2}}}(\mathbb{Z}) / \rho_{\mu, \frac{\varsigma}{\sqrt{2}}}(\mathbb{Z})}$ can be approximated with k bits of precision within $O(k^{3/2} \cdot \text{polylog}(k))$ time, which proves Lemma A.27.

A.6. Discrete Gaussians

Recall, for $n \in \mathbb{N}_{>0}$ and any parameter $s > 0$, we consider the n -dimensional *Gaussian function*

$$\rho_s^{(n)} : \mathbb{R}^n \rightarrow \mathbb{C}, \quad x \mapsto e^{-\frac{\pi\|x\|^2}{s^2}},$$

where we drop the (n) whenever it is clear from the context.

Lemma A.33. *We have*

$$|\rho_s(x) - \rho_s(y)| \leq \frac{\pi}{s^2} \cdot \|x - y\| \|x + y\| \cdot \rho_{2s}(x - y) \rho_{2s}(x + y).$$

Proof. We have, using the inequality $|1 - x| \leq |\ln(x)|$ (for all $x > 0$) and the reverse triangle inequality,

$$\begin{aligned} |\rho_s(x) - \rho_s(y)| &\leq \rho_s(x) |1 - \rho_s(x)/\rho_s(y)| \leq \frac{\pi}{s^2} \cdot \rho_s(x) |\|x\|^2 - \|y\|^2| \\ &\leq \frac{\pi}{s^2} \cdot \rho_s(x) \cdot \|x - y\| \|x + y\|. \end{aligned}$$

Since the bound above is symmetric in x and y , we might as well replace $\rho_s(x)$ by $\rho_s(y)$ in the rightmost expression, or even by their *harmonic*

mean $\sqrt{\rho_s(x)\rho_s(y)}$. Rewriting this harmonic mean $\sqrt{\rho_s(x)\rho_s(y)} = \rho_{2s}(x+y)\rho_{2s}(x-y)$ using multiplicative properties of the Gaussian function (see Lemma 2.23), we obtain the result. \square

Lemma A.34 (Bounds on the first and second moment of the discrete Gaussian). *Let $\Lambda \subseteq \mathbb{R}^n$ be a full-rank lattice and let $c \in \mathbb{R}^n$ and let $s > 4\sqrt{n} \cdot \lambda_n(\Lambda)$. Then, we have*

$$\begin{aligned} \frac{1}{\rho_s(\Lambda - c)} \sum_{\ell \in \Lambda} \rho_s(\ell - c) \|\ell - c\|^2 &\leq 2ns^2 \\ \frac{1}{\rho_s(\Lambda - c)} \sum_{\ell \in \Lambda} \rho_s(\ell - c) \|\ell - c\| &\leq 1 + 2ns^2. \end{aligned}$$

Proof. Using a result from Micciancio and Regev [MR07, Lm. 4.3] and the fact that $s > 4\sqrt{n}\lambda_n(\Lambda) > 2\eta_{1/2}(\Lambda)$, we directly obtain

$$\frac{1}{\rho_s(\Lambda - c)} \sum_{\ell \in \Lambda} \rho_s(\ell - c) \|\ell - c\|^2 \leq \left(\frac{1}{2\pi} + 1 \right) ns^2 \leq 2ns^2.$$

For the second bound, split up the sum in a part where $\|\ell - c\| \leq 1$ and $\|\ell - c\| > 1$. It is clear that the former must be bounded by 1, whereas the latter is bounded by $2ns^2$, by the fact that $\|\ell - c\| \leq \|\ell - c\|^2$ in that case. \square

Definition A.35. *Let $\mathbf{t} \in SL_m(\mathbb{R})$ be a diagonal matrix and let $\Lambda \subseteq \mathbb{R}^m$ be a full rank lattice. Then we define the distribution $\mathcal{G}_{\Lambda, s, \mathbf{t}, c}$ by the rule*

$$\mathcal{G}_{\Lambda, s, \mathbf{t}, c}(\ell) = \frac{\rho_s(\mathbf{t}(\ell - c))}{\rho_s(\mathbf{t}(\Lambda - c))}$$

Remark A.36. *Note that this definition coincides reasonably with the definition of the Gaussian distribution with a ‘variance matrix’ [Gut09, Ch. 5].*

Lemma A.37. *Let $\Lambda \subseteq \mathbb{R}^m$ be a full-rank lattice, $\varepsilon \in (0, \frac{1}{2})$, $c, \tilde{c} \in \mathbb{R}^m$, and $s \geq \eta_\varepsilon(\Lambda)$. Then*

$$\|\mathcal{G}_{\Lambda, s, c} - \mathcal{G}_{\Lambda, s, \tilde{c}}\| \leq 4\varepsilon + (\frac{2\pi}{s^2} + 4\pi n) \|c - \tilde{c}\|$$

Proof. By smoothing properties, we have $\rho_s(\Lambda - c), \rho_s(\Lambda - \tilde{c}) \in (1 - \varepsilon, 1 + \varepsilon)\rho_s(\Lambda)$. Allowing an extra error of 4ε , we can therefore replace the denominator in the definitions of $\mathcal{G}_{\Lambda, s, c}$ and $\mathcal{G}_{\Lambda, s, \tilde{c}}$ by $\rho_s(\Lambda)$.

$$\|\mathcal{G}_{\Lambda, s, c} - \mathcal{G}_{\Lambda, s, \tilde{c}}\| \leq 4\varepsilon + \frac{1}{\rho_s(\Lambda)} \sum_{\ell \in \Lambda} |\rho_s(\ell - c) - \rho_s(\ell - \tilde{c})|.$$

By Lemma A.33 (using the fact that $\rho_{s/2}(c - \tilde{c}) \leq 1$) and subsequently Lemma A.34, we have

$$\begin{aligned} \sum_{\ell \in \Lambda} |\rho_s(\ell - c) - \rho_s(\ell - \tilde{c})| &\leq \frac{\pi}{s^2} \|c - \tilde{c}\| \sum_{\ell \in \Lambda} \rho_{2s}(2\ell - (c + \tilde{c})) \|2\ell - (c + \tilde{c})\| \\ &\leq \frac{\pi}{s^2} (1 + 2ns^2) \|c - \tilde{c}\| \rho_s(\Lambda - \frac{c+\tilde{c}}{2}) \\ &\leq \frac{2\pi}{s^2} (1 + 2ns^2) \|c - \tilde{c}\| \rho_s(\Lambda). \end{aligned}$$

Combining the two bounds yields the result. \square

Lemma A.38. *Let $\Lambda \subseteq \mathbb{R}^m$ be a full-rank lattice, $c \in \mathbb{R}^m$, $\varepsilon, \delta \in (0, \frac{1}{2})$, $\mathbf{t} \in SL_m(\mathbb{R})$ be a diagonal matrix with² $|\mathbf{t} - 1| \leq \delta$. Additionally, assume that $s \geq \max(\eta_\varepsilon(\Lambda), \eta_\varepsilon(\mathbf{t}\Lambda))$. Then*

$$\|\mathcal{G}_{\Lambda, s, \mathbf{t}, c} - \mathcal{G}_{\Lambda, s, c}\| \leq 4\varepsilon + 2\pi n \delta$$

Proof. Since $\det(\mathbf{t}\Lambda) = \det(\Lambda) \prod_i \mathbf{t}_{ii} = \det(\Lambda)$, we have $\rho_s(\Lambda - c), \rho_s(\mathbf{t}(\Lambda - c)) \in (1 - \varepsilon, 1 + \varepsilon)\rho_s(\Lambda)$, by smoothing properties of the Gaussian function. Allowing an extra error of 4ε , we can therefore replace the denominator in the definitions of $\mathcal{G}_{\Lambda, s, c}$ and $\mathcal{G}_{\Lambda, s, \mathbf{t}, c}$ by $\rho_s(\Lambda)$.

$$\|\mathcal{G}_{\Lambda, s, \mathbf{t}, c} - \mathcal{G}_{\Lambda, s, c}\| \leq 4\varepsilon + \frac{1}{\rho_s(\Lambda)} \sum_{v \in \Lambda - c} |\rho_s(\mathbf{t}v) - \rho_s(v)|. \quad (\text{A.122})$$

²Here, we mean that the vector \mathbf{v} consisting of the diagonal elements of \mathbf{t} satisfies $|\mathbf{v} - 1| \leq \delta$ in the Euclidean norm.

By Lemma A.33, using the fact that $\rho_{2s}((\mathbf{t} - 1)v) \leq 1$ and $\|(\mathbf{t} - 1)v\| \leq \delta\|v\|$, we have

$$\begin{aligned} |\rho_s(\mathbf{t}v) - \rho_s(v)| &\leq \frac{\delta\pi}{s^2} \cdot \rho_{2s}((1 + \mathbf{t})v) \cdot \|v\| \cdot \|(\mathbf{t} - 1)v\| \\ &\leq \frac{\delta\pi}{s^2} \cdot \rho_s((1 + \mathbf{t})v) \cdot \|(\mathbf{t} - 1)v\|^2. \end{aligned} \quad (\text{A.123})$$

Where the last inequality follows from $\|v\| \leq \|(\mathbf{t} - 1)v\|$, which can be deduced by applying the triangle inequality on $\|v\|$ in the following way.

$$\begin{aligned} \|v\| &\leq \frac{1}{2}\|(1 + \mathbf{t})v\| + \frac{1}{2}\|(1 - \mathbf{t})v\| \leq \frac{1}{2}\|(1 + \mathbf{t})v\| + \frac{\delta}{2}\|v\| \\ &\leq \frac{1}{2}\|(1 + \mathbf{t})v\| + \frac{1}{2}\|v\|. \end{aligned}$$

Plugging Equation (A.123) into Equation (A.122), and applying Lemma A.34, we obtain

$$\|\mathcal{G}_{\Lambda,s/\mathbf{t},c} - \mathcal{G}_{\Lambda,s,c}\| \leq 4\varepsilon + \frac{\delta\pi}{s^2}(2ns^2) = 4\varepsilon + 2\pi n\delta.$$

□

Lemma A.39. *Let $\mathbf{t}_0, \mathbf{t}_1 \in SL_m(\mathbb{R})$ be diagonal matrices satisfying³ $|\mathbf{t}_0/\mathbf{t}_1 - 1| \leq \delta < 1/2$, let $\varepsilon \in (0, 1/2)$, let $c \in \mathbb{R}^m$ and let $\Lambda \subseteq \mathbb{R}^m$ be a full rank lattice. Let furthermore $s > \max(\eta_\varepsilon(\mathbf{t}_0\Lambda), \eta_\varepsilon(\mathbf{t}_1\Lambda))$.*

Then,

$$\|\mathcal{G}_{\Lambda_0,s,c} - \mathcal{G}_{\Lambda_0,s/\mathbf{t},c/\mathbf{t}}\| \leq 8\varepsilon + (2\pi n + (\frac{2\pi}{s^2} + 4\pi n)\|c\|) \cdot \delta.$$

Proof. We have, writing $\Lambda_0 = \mathbf{t}_0\Lambda$ and $\mathbf{t} = \mathbf{t}_1\mathbf{t}_0^{-1}$,

$$\begin{aligned} \sum_{\ell \in \Lambda} \left| \frac{\rho_s(\mathbf{t}_0\ell - c)}{\rho_s(\mathbf{t}_0\Lambda - c)} - \frac{\rho_s(\mathbf{t}_1\ell - c)}{\rho_s(\mathbf{t}_1\Lambda - c)} \right| &\leq \sum_{\ell_0 \in \Lambda_0} \left| \frac{\rho_s(\ell_0 - c)}{\rho_s(\Lambda_0 - c)} - \frac{\rho_s(\mathbf{t}\ell_0 - c)}{\rho_s(\mathbf{t}\Lambda_0 - c)} \right| \\ &= \|\mathcal{G}_{\Lambda_0,s,c} - \mathcal{G}_{\Lambda_0,s/\mathbf{t},c/\mathbf{t}}\| \leq \|\mathcal{G}_{\Lambda_0,s,c} - \mathcal{G}_{\Lambda_0,s,c/\mathbf{t}}\| + \|\mathcal{G}_{\Lambda_0,s,c/\mathbf{t}} - \mathcal{G}_{\Lambda_0,s/\mathbf{t},c/\mathbf{t}}\|. \end{aligned} \quad (\text{A.124})$$

³By this we mean that the vector $\mathbf{v} = \mathbf{t}_0/\mathbf{t}_1$ consisting of the diagonal elements of the matrix $\mathbf{t}_0/\mathbf{t}_1$ satisfies $|\mathbf{v} - 1| \leq \delta$ in the Euclidean norm.

Since $s \geq \eta_\varepsilon(\Lambda_0)$, by assumption, we have, by Lemma A.37,

$$\|\mathcal{G}_{\Lambda_0, s, c} - \mathcal{G}_{\Lambda_0, s, c/\mathbf{t}}\| \leq 4\varepsilon + \left(\frac{2\pi}{s^2} + 4\pi n\right)\|c - c/\mathbf{t}\| \leq 4\varepsilon + \left(\frac{2\pi}{s^2} + 4\pi n\right)\|c\| \cdot \delta,$$

since $\|1 - 1/\mathbf{t}\| \leq \|1 - \mathbf{t}_0/\mathbf{t}_1\| \leq \delta$ by assumption. Also, since $s \geq \eta_\varepsilon(\mathbf{t}\Lambda_0)$ (note that $\mathbf{t}\Lambda_0 = \mathbf{t}_1\Lambda$), we have, by Lemma A.38,

$$\|\mathcal{G}_{\Lambda_0, s, c/\mathbf{t}} - \mathcal{G}_{\Lambda_0, s/\mathbf{t}, c/\mathbf{t}}\| \leq 4\varepsilon + 2\pi n\delta.$$

Combining the bounds into Equation (A.124), we obtain the final claim. \square

Summary

The main topic of this PhD thesis is the *Arakelov ray class group* of a number field, an algebraic object that contains both the ideal class group structure and the unit group structure. The main result consists of the fact that certain specific *random walks* on the Arakelov ray class group result in a target point that is uniformly distributed on this group, under the assumption of an extended version of the Riemann Hypothesis (Chapter 4). Almost all other results of this work are consequences of this fact.

As a first direct application, using these random walks on the Arakelov class group one can show that finding a short vector in a *arbitrarily chosen* ideal lattice is no harder than finding a short vector in a *random ideal lattice* of a fixed number field (Chapter 5). In other words, finding short vectors in the ‘most difficult’ ideal lattice is not much harder than finding short vectors in a random ideal lattice.

A second application uses these random walks to rigorously and efficiently sample elements from ideals of number fields, in such a way that the quotient of this element and the ideal lies in a pre-chosen ideal set. The success probability of this sampling procedure turns out to be proportional to the analytic number-theoretic *density* of this ideal set. One obtains a particularly interesting application of this result when one chooses the ideal set to be the set of prime ideals and when the number field equals a cyclotomic number field. In that case the theorem reads: one can efficiently sample elements in ideals of cyclotomic number fields such that the quotient of the concerning element and ideal is near-prime, i.e., a product of a large prime ideal and several very small prime ideals (Chapter 6).

The purpose of the above sampling algorithm is to transform heuristic arguments into rigorous proofs in certain number-theoretic algorithms, like ideal class group and unit group algorithms. We successfully achieve this goal for an algorithm that computes the power residue symbol: we give a formal proof of the polynomial running time of that algorithm. Before the writing of this thesis, this running time was only heuristically estimated to be polynomially bounded (Chapter 7).

A more self-contained part of this thesis consists of a quantum algorithm of the *continuous hidden subgroup problem* and a full, rigorous analysis thereof (Chapter 3). This algorithm can be applied to *compute* Arakelov ray class groups explicitly; though this is still a topic of research.

Samenvatting

Het hoofdonderwerp van deze dissertatie is de *Arakelov straalklassegroep* van een getallenlichaam, een algebraïsch object dat zowel de ideaalklassesstructuur als de eenhedenstructuur van een getallenlichaam omvat. Het hoofdresultaat van deze thesis betreft het feit dat zekere specifieke *toevalsbewegingen* op de Arakelov straalklassegroep resulteren in een eindpunt dat uniform random verdeeld is over deze groep, onder aanname van een uitgebreide variant van de Riemann hypothese (hoofdstuk 4). Bijna alle andere resultaten in dit werk vloeien voort uit dit feit.

Als een direct gevolg, kunnen we met behulp van deze toevalsbewegingen op de Arakelov klassegroep aantonen dat het vinden van een kortste vector in een gegeven (mogelijk ingewikkeld) ideaalrooster computationally niet veel moeilijker is dan het vinden van een kortste vector in een random ideaalrooster van een vast gekozen getallenlichaam (hoofdstuk 5).

Een ander resultaat maakt gebruik van deze toevalsbewegingen om rigoureus en op een efficiënte wijze elementen uit idealen van een getallenlichaam te samplen op zodanige manier dat het quotiënt van het betreffende element en ideaal uit een bepaalde vooraf gekozen ideaalverzameling komt. De succeskans van dit samplen is dan evenredig met de analytisch-getaltheoretische *dichtheid* van de ideaalverzameling op een zeker punt. Een interessante toepassing van dit resultaat verkrijgt men wanneer men voor de betreffende ideaalverzameling de verzameling van priemidealen neemt en voor het getallenlichaam een cyclotomisch lichaam. In dat geval luidt de stelling dat men op efficiënte wijze elementen in idealen van cyclotomische lichamen kan vinden zodanig dat het quotiënt van het betreffende element en ideaal

‘bijna’ een priemideaal is. Dat wil zeggen dat dit quotiënt bestaat uit het product van een groot priemideaal en eventueel meerdere kleinere priemidealen (hoofdstuk 6).

Het doel van bovenstaand sampling algoritme is om heuristische argumenten voor bepaalde algoritmen in de getaltheorie, zoals ideaalklasse- en eenheden-groepalgoritmes, te vervangen door rigoureuze bewijzen. In deze dissertatie is het op deze manier gelukt om formeel te bewijzen dat een algoritme voor het machtrestsymbool, een polynomiale looptijd heeft. Van de looptijd van dit algoritme was voorheen alleen *heuristisch* aangetoond dat deze polynomiaal begrensd was (hoofdstuk 7).

Een deel van deze thesis dat wat meer op zichzelf staat, betreft een kwantumalgoritme voor een continue variant van het *hidden subgroup problem* en een volledige analyse van dit algoritme (hoofdstuk 3). Dit algoritme kan toegepast worden om Arakelov straalklassen-groepen expliciet te berekenen; dit moet echter nog wel nauwkeurig onderzocht en bewezen worden.

Acknowledgments

I would like to thank my advisors Prof. dr. Léo Ducas, dr. Benjamin Wesolowski and Prof. dr. Ronald Cramer for their expert supervision.

Léo, my promotor, has shown unwavering support during my research and during the writing of this thesis as well. He always enthusiastically shared his knowledge and insights with me throughout this time. Much of the work in this dissertation is the result of our shared dedication to rigorous run-time analyses of algorithms in cryptography.

Benjamin, my co-promotor, greatly increased my interest in analytic number theory and its usage in cryptography by generously sharing his expertise. I am grateful for our cooperation, which was marked by ample room for discussing mathematical intricacies and having a good laugh as well.

Ronald, my second promotor, is gratefully acknowledged for his guidance and constructive criticism during the preparation of the final version of this thesis. I am very grateful to him for putting me on the right track scientifically and non-scientifically during the more difficult times of my PhD track.

I am also very grateful to the members of the Doctorate Committee for reading my thesis and providing useful feedback.

Additionally, I would like to thank my colleagues – from the CWI and Universiteit Leiden, but also from abroad – for providing me an inspiring and positive environment.

To my family, friends and partner I would like to say: Jullie zijn een onvervangbare en onmisbare steun geweest tijdens mijn promotietraject. Enorm

Acknowledgments

bedankt voor jullie geduld, begrip en liefde — op de betere, maar ook vooral op de zwaardere momenten van deze periode.

Curriculum Vitae



Koen de Boer was born in Nijmegen, the Netherlands, on August 22, 1991. He grew up in the city Oss in Brabant, where he obtained his high school diploma from Titus Brandsma Lyceum in 2008.

After deciding to study Mathematics at the Radboud Universiteit of Nijmegen, he obtained his bachelor degree *cum laude* in 2012 and his master's degree *summa cum laude* in 2016. His master thesis, "Computing the Power Residue Symbol", was written under supervision of dr. W. Bosma and Prof. dr. H.W. Lenstra.

In 2016, Koen obtained a PhD position at the Universiteit Leiden under supervision of Prof. dr. L. Ducas, Prof. dr. R. Cramer and dr. B. Wesolowski, to do research in the Cryptology Group at Centrum Wiskunde & Informatica (CWI) in Amsterdam.

In 2022, he started to work as a post-doc at the Universiteit Leiden.

Publications

- K. de Boer and C. Pagano (2017). "Calculating the Power Residue Symbol and Ibeta: Applications of Computing the Group Structure of the Principal Units of a p-adic Number Field Completion." In: *Proceedings of the 2017 ACM on International Symposium on Symbolic and Algebraic Computation (ISSAC '17)*. Association for Computing Machinery, New York, NY, USA.

- K. de Boer, L. Ducas, S. Jeffery, R. de Wolf (2018). “Attacks on the AJPS Mersenne-Based Cryptosystem.” In: Post-Quantum Cryptography. PQCrypto 2018. *Lecture Notes in Computer Science*, vol 10786. Springer, Cham.
- K. de Boer, L. Ducas, S. Fehr (2020). “On the Quantum Complexity of the Continuous Hidden Subgroup Problem.” In: Advances in Cryptology – EUROCRYPT 2020. *Lecture Notes in Computer Science*, vol 12106. Springer, Cham.
- K. de Boer, L. Ducas, A. Pellet-Mary, B. Wesolowski (2020). “Random Self-reducibility of Ideal-SVP via Arakelov Random Walks.” In: Advances in Cryptology – CRYPTO 2020. *Lecture Notes in Computer Science*, vol 12171. Springer, Cham.

List of Symbols

$\ \cdot\ _{p,G}$	The Haar-measure induced p -norm on functions $G \rightarrow \mathbb{C}$, where G is a locally compact abelian group G . In this thesis, p is either 1, 2 or ∞ in this context. The subscript G is often suppressed, as well as the subscript p in the case of $p = 2$ (page 43)
$f _H^H$	The periodization of a function $f : G \rightarrow \mathbb{C}$ with respect to a subgroup $H \subseteq G$ (page 45)
$f _H$	The restriction of a function $f : G \rightarrow \mathbb{C}$ with respect to a subgroup $H \subseteq G$ (page 45)
$ \cdot\rangle, \langle\cdot , \langle\cdot \cdot\rangle$	The ket, bra and bra-ket notation, used for quantum states in a quantum Hilbert space \mathcal{H} (page 41)
$\lfloor \cdot \rfloor, \lfloor \cdot \rfloor, \lceil \cdot \rceil$	Respectively, rounding to the nearest integer ($x \in [-\frac{1}{2}, \frac{1}{2}]$ rounds to 0), rounding down and rounding up
\star	The convolution operation on functions on a locally abelian group G (page 45)
$\langle \cdot \rangle$	The diagonal embedding of K into the Arakelov divisor group Div_K . The notation $\langle \mathfrak{p} \rangle$ and $\langle \nu \rangle$ for prime ideals and places is also used for the generators in the Arakelov divisor group (page 59)
$\hat{\cdot}$	The dual group of a locally compact abelian group, e.g., \hat{G} is the dual group of G (page 42)
$\tilde{\cdot}$	An approximation; for example, \tilde{B} indicates an approximation of B
\cdot^0	The subgroup of elements of norm or degree one, where the norm or degree is induced by the associated number field. For example, Div_K^0 , Pic_K^0 (page 60), $K_{\mathbb{R}}^0$ (page 53), \mathcal{J}_K^0 , \mathcal{C}_K^0 (page 148)
$\cdot^{\mathfrak{m}}$	The ray analogue of a number field related group, involving \mathfrak{m} . For example, $\text{Div}_{K^{\mathfrak{m}}}$, $\text{Pic}_{K^{\mathfrak{m}}}$ (page 59), $\mathcal{I}_K^{\mathfrak{m}}$ (page 54), $\text{Cl}_K^{\mathfrak{m}}$ (page 62)
$\ddot{\cdot}$	A discretized analogue of a continuous object. For example, $\ddot{\mathcal{D}}$ for a discretized version of a continuous distribution

$\left(\frac{\alpha}{\mathfrak{b}}\right)_{m,K}$	The m -th power residue symbol, where the top argument is an element of K and the bottom argument is an ideal of a number ring of K ; if K is clear from context, this notation is dropped (page 238)
$(\alpha, \beta)_{\mathfrak{p}}$	The m -th Hilbert symbol, where both arguments α, β are elements of the \mathfrak{p} -adic completion $K_{\mathfrak{p}}$ of the associated number field K ; the ‘power’ m is always clear from the context and not included in the notation (page 259)
(ϕ, M)	The signature symbol of the automorphism ϕ on a finite admissible module M (page 245)
$(x_{\sigma})_{\sigma}, (r_{\sigma})_{\sigma}$	Elements in $K_{\mathbb{R}}$, written as vectors indexed by the embeddings of K into \mathbb{C} (page 53)
$\mathbf{1}_G$	The unit character on the locally compact abelian group G
$\mathfrak{a}, \mathfrak{b}, \mathfrak{c}, \dots$	Ideals of the ring of integers of a number field, elements of \mathcal{I}_K (page 54)
$[\mathfrak{a}], [\mathfrak{b}], [\mathfrak{c}], \dots$	Ideal classes, elements of Cl_K (page 55)
$\mathbf{a}, \mathbf{b}, \mathbf{c}, \dots$	Arakelov divisors, elements of Div_K (page 59)
$[\mathbf{a}], [\mathbf{b}], [\mathbf{c}], \dots$	Arakelov classes, elements of Pic_K (page 60)
$\mathbf{a}_f, \mathbf{a}_{\infty}$	The finite part and respectively infinite part of an Arakelov divisor \mathbf{a} (page 61)
$\mathcal{B}_r(x)$	The ball of radius r around x with respect to the 2-norm (page 77)
$\mathcal{B}_r(X)$	The union of balls $\bigcup_{x \in X} \mathcal{B}_r(x)$ over all $x \in X$ (page 91)
$r\mathcal{B}_{\infty}, (r_{\sigma})_{\sigma}\mathcal{B}_{\infty}$	The box of radius r around 0 with respect to the ∞ -norm (page 210) respectively the distorted box of component-wise radius $(r_{\sigma})_{\sigma}$ (page 226)
C	In Chapter 3, the ‘target set’ of the algorithm, i.e., the set of good outcomes of a measurement (page 102)
$C(r, \mathcal{N}(\mathfrak{c}))$	The volume of the simplex with edge length $(n \log r - \log \mathcal{N}(\mathfrak{c}))$ and dimension \mathfrak{r} , which equals $(n \log r - \log \mathcal{N}(\mathfrak{c}))^{\mathfrak{r}} / \mathfrak{r}!$ whenever $\mathcal{N}(\mathfrak{c}) \leq r^n$ and zero otherwise (pages 215 and 279)
\mathcal{C}_K	The idèle class group of a number field (page 55)
Cl_K	The class group of a number field (page 55)
Cl_K^{m}	The ray class group of a number field (page 62)
\mathcal{C}_M	The hypercircle $\{(x_{\sigma})_{\sigma} \in K_{\mathbb{R}} \mid x_{\sigma} = M\}$ (page 175)
$\tilde{\mathcal{C}}_M$	The discretized hypercircle (page 192)
$\text{cov}_2(\Lambda), \text{cov}_{\infty}(\Lambda)$	The covering radius of a lattice Λ with respect to the 2-norm or the ∞ -norm respectively (page 72)
d	The map sending \mathcal{I}_K to Div_K by using the valuations of the prime ideals as coefficients for the finite places of the Arakelov divisor (page 61)

d^0	The map sending \mathcal{I}_K to Div_K^0 by using the valuations of the prime ideals as coefficients for the finite places of the Arakelov divisor, and a fraction of the negative logarithmic norm of the ideal at the infinite places (page 61)
\mathcal{D}	Generally, a distribution. In Chapter 6, it is a distribution over $\text{Div}_{K^m}^0$
$[\mathcal{D}]$	The $K^{m,1}$ -periodization $\mathcal{D} ^{K^{m,1}}$ of a distribution \mathcal{D} over $\text{Div}_{K^m}^0$ (page 210)
$\mathcal{D}_{x\mathfrak{a}}$	The distribution representation of an ideal lattice $x\mathfrak{a}$ (page 174)
\mathbb{D}^m	The group $\frac{1}{q}\mathbb{Z}^m/\mathbb{Z}^m \subseteq \mathbb{T}^m$, a q -discretized version of the unit torus (page 42)
$\mathbb{D}_{\text{rep}}^m$	The standard representation $\frac{1}{q}\mathbb{Z}^m \cap [-\frac{1}{2}, \frac{1}{2})^m$ of \mathbb{D}^m (page 42)
$\hat{\mathbb{D}}^m$	The dual of \mathbb{D}^m , isomorphic to $\mathbb{Z}^m/q\mathbb{Z}^m$ (page 42)
$\hat{\mathbb{D}}_{\text{rep}}^m$	The standard representation $\mathbb{Z}^m \cap [-\frac{q}{2}, \frac{q}{2})^m$ of $\hat{\mathbb{D}}^m$ (page 42)
$\deg(\cdot)$	The degree of an Arakelov divisor; a weighted sum of the coefficients associated with the places. Equivalently, the logarithm of the determinant of the ideal lattice associated with the Arakelov divisor (page 60)
$\det(\cdot)$	The determinant of a matrix, or, the determinant of a lattice Λ , which is equal to its covolume $\text{Vol}(\Lambda)$ (page 72)
Div_K	The Arakelov divisor group of a number field K , consisting of formal sums of places of K (page 59)
Div_{K^m}	The subgroup of the Arakelov divisor group consisting of formal sums not involving the places dividing the modulus \mathfrak{m} (page 59)
$\text{Exp}(\mathbf{a})$	The exponentiation map sending an Arakelov divisor $\mathbf{a} \in \text{Div}_K$ to an ideal lattice in IdLat_K (page 73)
$\text{Exp}(\mathbf{a})_\tau^\times$	The τ -equivalent generators of the Arakelov ray divisor \mathbf{a} (page 208)
\mathbf{f}	In Chapter 3, the periodic function over \mathbb{R}^m that ‘hides’ the lattice Λ (page 81)
$\mathcal{F}_G\{\cdot\}$	The Fourier transform with respect to the locally compact abelian group G (page 44)
$\mathcal{G}_{X,s}$	The (discrete) Gaussian distribution with deviation s , where the structure of the space X determines whether \mathcal{G} is continuous or discrete (page 79)
\mathbf{h}	In Chapter 3, the ‘wave packet variant’ of the periodic function over \mathbb{R}^m that hides the lattice Λ (page 102)
H	The hyperplane $\text{Log}(K_{\mathbb{R}}^0)$ in $\text{Log}(K_{\mathbb{R}})$ where the Logarithmic unit lattice $\Lambda_K = \text{Log}(\mathcal{O}_K^\times)$ lives in (page 54). Occasionally, a subgroup of a locally abelian group G
\mathcal{H}	In Chapter 3, a finite-dimensional quantum Hilbert space (page 41). In Chapter 4, a Hecke operator (page 142).

$\mathcal{H}_{\mathcal{P}}$	The Hecke operator with respect to a finite set of prime ideals \mathcal{P} ; this set is omitted in the notation if it is clear from context (page 142)
h_K	The class number $ \text{Cl}_K $ (page 53)
h_K^+	The class number of the maximal totally real subfield of K (page 55)
\mathcal{I}_K	The group of fractional ideals of the ring of integers of a number field K (page 54)
$\mathcal{I}_K^{\mathfrak{m}}$	The subgroup of \mathcal{I}_K consisting of ideals coprime to a modulus \mathfrak{m} (page 54)
\mathcal{J}_K	The idèle group of the number field K (page 55)
IdLat_K	The group of ideal lattices of the number field K (page 73)
K	A finite-degree number field (page 53)
$K^{\mathfrak{m},1}$	The multiplicative subgroup of K generated by the elements in \mathcal{O}_K that are equivalent to 1 modulo the modulus \mathfrak{m} (page 55)
$K^{\mathfrak{m}}$	The multiplicative subgroup of K generated by the elements in \mathcal{O}_K that are invertible modulo the modulus \mathfrak{m} (page 55)
$K_{\mathbb{R}}$	The tensor product $\mathbb{R} \otimes_{\mathbb{Z}} K$ where K is a number field; also, co-domain of the Minkowski embedding (page 53)
$K_{\mathbb{R}}^0$	The subgroup of $K_{\mathbb{R}}$ consisting of those elements whose K -induced algebraic norm equals 1 (page 54)
$K_{\nu}, K_{\mathfrak{p}}$	The completion of K with respect to the place ν or prime \mathfrak{p} (page 55)
ℓ	Generally, a lattice point $\ell \in \Lambda$
ℓ^*	Generally, a dual lattice point $\ell^* \in \Lambda^*$
$L(\chi, s)$	The L-function associated with a Hecke character χ of a number field K (page 56)
$L_p(G)$	The metric vector space of measurable functions $f : G \rightarrow \mathbb{C}$ on a locally abelian group G for which the p -norm $\ f\ _{p,G}$ is well-defined and finite (modulo functions with norm zero) (page 43)
$\text{Lip}(f)$	The Lipschitz constant of a function f between two normed spaces (page 90)
Log	The logarithmic map $K \rightarrow \text{Log}(K_{\mathbb{R}})$ defined by taking the logarithm of the absolute value of each component of the Minkowski embedding (page 54)
m	In Chapter 3, the dimension of the hidden lattice; in Chapter 7 the m in the m -th power residue symbol
\mathfrak{m}	An ideal modulus of a number field, consisting of a formal product of finite places of that number field (page 54)
n	The degree $[K : \mathbb{Q}]$ of the number field K (page 53)
$n_{\mathbb{R}}$	The number of real embeddings $K \hookrightarrow \mathbb{R}$ (page 53)

$n_{\mathbb{C}}$	The number of conjugate pairs of complex embeddings $K \hookrightarrow \mathbb{C}$ (page 53)
$\mathcal{N}(\cdot)$	The algebraic norm of a number field element or ideal (page 55)
$o(\cdot)$	The Bachmann-Landau Small-o notation
$O(\cdot)$	The Bachmann-Landau Big-O notation
$\tilde{O}(\cdot)$	The soft-O notation, ignoring polylogarithmic factors
\mathcal{O}_K	The ring of integers of the number field K (page 53)
\mathcal{O}_K^\times	The unit group of the number field K (page 54)
$\mathcal{O}_{K^{\mathfrak{m},1}}^\times$	The ray unit group of the number field K with respect to the modulus \mathfrak{m} , i.e., $\mathcal{O}_K^\times \cap K^{\mathfrak{m},1}$ (page 62)
$\text{ord}_{\mathfrak{p}}$	The valuation with respect to the prime ideal \mathfrak{p} (page 54)
\mathfrak{p}	A prime ideal of a number field (page 54)
\mathfrak{p}_ν	A prime ideal of a number field uniquely associated with the finite place ν (page 53)
Princ_K	The subgroup of principal ideals in \mathcal{I}_K , i.e., those generated by a single element in K (page 55)
\mathfrak{q}	A prime ideal of a number field (page 54)
$\mathfrak{q}_\infty(\chi)$	The infinite part of the analytic conductor of a Hecke character $\chi \in \widehat{\text{Pic}}_{K^{\mathfrak{m}}}^0$ (page 151)
q	The discretization parameter in the continuous hidden subgroup quantum algorithm (page 100)
Q	$\log(q)$, the numbers of qubits ‘per dimension’ in the continuous hidden subgroup quantum algorithm (page 87)
r	Generally, the radius of a ball or box in a vector space; in Chapter 3, part of the definition of a function being (r, ϵ) -separating (page 91)
\mathfrak{r}	The rank of the unit group of a number field K , which equals $n_{\mathbb{R}} + n_{\mathbb{C}} - 1$ (page 54)
R_K	The regulator of the number field K , strongly related to the volume of T (page 53)
\mathcal{S}	In Chapter 3, the space of quantum states (page 90). In Chapters 6 and 7, a set of integral ideals of the ring of integers of a number field K (page 209)
$\mathcal{S}^{\mathfrak{m}}$	A set of integral ideals of the ring of integers of a number field K that are coprime with the modulus \mathfrak{m} (page 209)
\mathcal{S}_B	The set of all B -smooth integral ideals of \mathcal{O}_K , i.e., all integral ideals having only prime ideal factors with norm $\leq B$ (page 205)
$ \mathcal{S}(t) $	The number of ideals in \mathcal{S} with norm bounded by t (page 209)
s	The deviation for the Gaussian function or the (discrete) Gaussian distribution. Occasionally, input variable of zeta functions and L-functions

$\text{span}(\cdot)$	The linear subspace spanned by the vectors or the lattice within the brackets
\mathbb{T}^m	The unit torus $\mathbb{R}^m / \mathbb{Z}^m$ (page 42)
$\mathbb{T}_{\text{rep}}^m$	The standard representation $[-\frac{1}{2}, \frac{1}{2})^m$ of the unit torus \mathbb{T}^m (page 42)
T	The logarithmic unit torus $H / \text{Log}(\mathcal{O}_K^\times)$ (page 54)
T^m	The logarithmic ray unit torus $H / \text{Log}(\mathcal{O}_{K^m,1}^\times)$ (page 62)
$\mathcal{U}(X)$	The uniform distribution over the compact space X (page 65)
$\text{Vol}(\cdot)$	Volume of the compact abelian group with respect to the fixed given Haar measure (page 42), or, the covolume of a lattice (also called the determinant of the lattice, (page 72))
$\mathcal{W}_{\text{Pic}_{K^m}^0}(B, N, s)$	The random walk distribution over the Arakelov ray class group $\text{Pic}_{K^m}^0$ with prime ideal norm bound B , number of steps N and Gaussian deviation s (page 140)
$x\mathfrak{a}, y\mathfrak{b}$	Ideal lattices, elements of IdLat_K (page 73)
\mathbb{Z}_H	Orthogonal discretization of the hyperplane H where the log unit lattice $\Lambda_K = \text{Log}(\mathcal{O}_K^\times)$ lives in (page 191)
$\alpha, \beta, \gamma, \dots$	Generally, elements of a number field K (page 53)
β_z	Banaszczyk's function $z \mapsto \left(\frac{2\pi e z^2}{n}\right)^{n/2} e^{-\pi z^2}$ (page 77)
Γ_K	The maximum of the quotient between the outermost successive minima $\lambda_n(x\mathfrak{a})/\lambda_1(x\mathfrak{a})$ over all ideal lattices $x\mathfrak{a} \in \text{IdLat}_K$ for a fixed number field K (page 75)
δ	In Chapter 3, the relative distance error in the dual sampling algorithm (page 93). In the rest of the thesis, generally a small distance or error
$\delta_{\mathcal{S}}[x]$	The local density of the ideal set \mathcal{S} around norm x (page 209)
Δ_K	The discriminant of the number field K (page 53)
ε	A small error parameter in $[0, 1]$, often indicating the failure probability of an algorithm
ϵ	A parameter in the definition of a function being (r, ϵ) -separating in Chapter 3 (page 91)
$\zeta(s)$	The Riemann zeta function
ζ_m	A primitive m -th root of unity
$\zeta_K(s)$	The Dedekind zeta function with respect to the number field K (page 56)
η	In the dual lattice sampling algorithm of Chapter 3, the failure probability of the algorithm (page 93). In the rest of the thesis, a small error or sometimes a unit $\eta \in \mathcal{O}_K^\times$
$\eta_\varepsilon(\Lambda)$	The smoothing parameter, the smallest $s > 0$ such that $\rho_{1/s}(\Lambda^* \setminus \{0\}) \leq \varepsilon$ (page 77)
λ_1^*	The first successive minimum $\lambda_1(\Lambda^*)$ of the dual lattice, whenever the lattice Λ is clear from context (page 72)

$\lambda_j(\Lambda)$	The j -th successive minimum of the lattice Λ with respect to the 2-norm (page 72)
$\lambda_j^{(\infty)}(\Lambda)$	The j -th successive minimum of the lattice Λ with respect to the ∞ -norm (page 72)
λ_χ	The eigenvalue of the character $\chi \in \widehat{\text{Pic}_{K^m}^0}$ under the Hecke operator \mathcal{H}
Λ	A lattice, i.e., a discrete subgroup of a Euclidean vector space (page 72)
Λ^*	The dual of the lattice Λ (page 72)
Λ_K	The log unit lattice $\text{Log}(\mathcal{O}_K^\times)$
Λ_{K^m}	The ray log unit lattice $\text{Log}(\mathcal{O}_{K^{m,1}}^\times) = \text{Log}(\mathcal{O}_K^\times \cap K^{m,1})$
μ_K	The group of roots of unity of the number field K (page 53)
ν	A formal place, associated with an absolute value $ \cdot : K \rightarrow \mathbb{R}_{>0}$ on a number field K (page 53)
ν_σ	The place associated with the absolute value induced by the embedding $\sigma : K \rightarrow \mathbb{C}$ (page 53)
ρ_s	The Gaussian function $x \mapsto e^{-\pi\ x\ ^2/s^2}$ (page 76)
ρ_K	The residue $\lim_{s \rightarrow 1} (1-s)\zeta_K(s)$ of the Dedekind zeta function at $s = 1$ (page 56)
σ	An embedding from a number field K into the complex numbers \mathbb{C} (page 53)
ς	The deviation for the Gaussian function or the (discrete) Gaussian distribution whenever s is already used
$\phi(m)$	The Euler indicator function, $\phi(m) = (\mathbb{Z}/m)^* $ for $m \in \mathbb{N}_{>0}$
$\phi(\mathfrak{m})$	The generalized Euler indicator function for ideals $\mathfrak{m} \subseteq \mathcal{O}_K$, $\phi(\mathfrak{m}) = (\mathcal{O}_K/\mathfrak{m})^* $ (page 62)
χ	A character $\chi \in \hat{G}$ of a locally compact abelian group G , i.e., a continuous group homomorphism from G to \mathbb{C} (page 42)

Index

- analytic conductor, 151
 - bound on the, 152
 - informal description of the, 139
- Arakelov ray class groups, 60
 - earlier (cryptographic) work related to, 133
 - informal description of, 129
 - motivation for studying, 132
 - example of, 66
 - volume of, 64
- Arakelov ray divisor, 59
 - (τ -equivalent) generator of an, 208
 - finite and infinite part of a, 61
- Artin symbols
 - algorithm to compute, 261
 - Dedekind zeta function and its influence on computing, 263
- Banaszczyk's bound, *see* Gaussian distribution
- class group and unit group of a number field, 53
 - quantumly computing the, 90
 - logarithmic unit lattice, *see* logarithmic unit lattice
- concentrated
 - (R, q) -concentrated lattice distribution, 117
 - continuous hidden subgroup problem, 85
 - an informal description of the, 81
 - quantum algorithm solving the, 97
 - correctness of the, 92
 - summary of the, 87
 - research directions relating to the, 88
 - covering radius, *see* lattice
 - covolume, *see* lattice
 - cyclotomic units, 188
 - relation between the random walk theorem and the, 160

- Dedekind zeta function, 55
 - bound on the residue of the, 271
 - influence on the computation of Artin symbols of the, *see* Artin symbols
- determinant, *see* lattice
- discretization
 - errors in the continuous HSP quantum algorithm caused by, 100
 - need for discretization in the reduction algorithm, *see* worst-case to average-case reduction on ideal lattices
- distribution
 - average-case distribution for ideal lattices, 180
 - Gaussian, *see* Gaussian distribution
- dual lattice of the hidden lattice
 - recovering the full, 94, 122
- dual lattice sampling problem
 - an informal description of the, 87
 - analysis of the quantum algorithm solving the, 107
 - definition of the, 93
 - quantum algorithm solving the, 101
 - short analysis of the quantum algorithm solving the, 104
 - theorem about the quantum algorithm solving the, 93, 115
- evenly distributed
 - p -evenly distributed lattice distribution, 117
- exact sequences
 - kernel-cokernel, 280
 - the Arakelov ray class group within a diagram of, 61
- extended Riemann hypothesis, 55
- Fourier analysis
 - on the Arakelov ray class group (informal), 131
 - on the ray unit torus, 152
 - on locally compact abelian groups, 43
 - on the Arakelov ray class group, 65, 155
- Gaussian distribution
 - discrete, 296
 - notation for the discrete and continuous, 79
 - results on shifting a discrete, 298
 - tail bounds on the discrete, 76
- Gaussian quantum state
 - setting up the initial, 96, 294
- generator
 - of an Arakelov ray divisor, *see* Arakelov ray divisor

Hecke operator, 142
 bounds on the eigenvalues of the, 143
 informal description of the, 138

hidden lattice
 recovering the basis of the, 95, 126

hidden lattice problem, *see* continuous hidden subgroup problem

Hilbert symbols, 260
 using power residue symbols to compute, 259

HSP-oracle
 (a, r, ϵ) -HSP oracle, 91

ideal density, *see* local density of an ideal set

ideal lattices, 73
 associated to a Arakelov divisor, 73
 definition of, 172
 distribution representation of, 174
 algorithm for the, 175
 definition of the, 175
 discrete algorithm for the, 194
 properties of, 176
 invariants of, 75
 isometry of, 73
 modulo isometry, *see* Arakelov ray class groups
 need for an efficient representation of, 173
 the Arakelov class group and its relation to, 74

ideals
 sampling, *see* sampling
 set of, *see* local density of an ideal set
 smooth, 228

lattice
 invariants of a, 71

Lipschitz continuity, 90
 influence on the Fourier coefficients, 79

local density of an ideal set, 210
 main theorem relating the sampling probability to the, 210
 proof of the, 216

logarithmic unit lattice, 53
 volume of the, 268

number field
 invariants of a, 53
 norm on a, 72

- period finding, *see* continuous hidden subgroup problem
- periodization and restriction, 45
- Poisson summation formula, 46
- power residue symbol
 - algorithm computing the
 - difference with an earlier heuristic algorithm, 237
 - earlier work, 236
 - efficiency of the, 259
 - cyclotomic, *see* power residue symbol in cyclotomic fields
 - power residue symbol in cyclotomic fields
 - algorithm computing the, 257
 - correctness of the, 257
 - informal description of the, 255
 - role of the random walk in the, 256
- probability-density correspondence, *see* local density of an ideal set
- random walk distribution on the Arakelov ray class group, 140
 - an informal description of the, 130
 - intuitive argument for the uniformity of, 136
 - algorithm mimicking the, 229
 - correctness of the, 231
- random walk theorem for the Arakelov ray class group, 141, 161
 - applications of the, 162
 - interpretation of the, 161
 - proof overview of the, 138
- reduction algorithm
 - for power residue symbols, *see* power residue symbol
 - on ideal lattices, *see* worst-case to average-case reduction on ideal lattices
- representation
 - of ideal lattices, *see* ideal lattices
- Riemann hypothesis, *see* extended Riemann hypothesis
- rigorously sampling elements in ideals, 204
 - applications of, 206
 - earlier work related to, 207
 - the role of the random walk theorem in, 205
- sampling
 - ideal sampling algorithm, 229
 - correctness of the, 228
 - uniformly sampling an element in a box, 226
- sampling probability of ideals, *see* local density of an ideal set
- separating
 - (r, ϵ) -separating function, 91

shortest vector problem
Hermite variant of the, 72
self-reduction of the, *see* worst-case to average-case reduction on ideal lattices
simplex
volume of the, 267
smoothing parameter, *see* gaussian distribution
successive minimum, *see* lattice

trigonometric approximation
usage in Fourier analysis, 51
Yudin's result, 283

unit group, *see* class group and unit group of a number field

worst-case to average-case reduction
informal description, 165
other works using random walks to obtain a, 170
on ideal lattices, *see* worst-case to average-case reduction on ideal lattices
worst-case to average-case reduction on ideal lattices
earlier works on, 170
informal description, 168
relation between cryptography and the, 170
algorithm for the, 181
closeness of discrete and continuous, 197
correctness of the, 183
discrete version of the, 196
explanation of the, 180
discretization of, 189
need for, 188
main theorem, 184
informal version, 167
loss of shortness quality in the, 186