S-unit attacks on structured lattice-based cryptosystems

Henry Bambury
St John's College
University of Oxford

A dissertation submitted for the degree
MSc Mathematics and Foundations of Computer Science

Trinity 2022

Acknowledgements

I would like to thank Christophe Petit for his supervision on this dissertation, as well as my college advisor Stefan Kiefer for his support throughout the year.

I would also like to thank my college for giving me the opportunity to attend the *Summer School in Post-Quantum Cryptography* organised in Budapest, while I was studying for this dissertation.

Finally, I am grateful to the *Direction Générale de l'Armement* for generously funding my degree.

Abstract

Amidst concerns that society will need new forms of encryption to protect itself against the arrival of the (in)famous quantum computer, alternatives to integer factoring and discrete logarithm problems have started to rise. The most promising amongst them relies on the problem of finding short vectors in lattices endowed with a particular number theoretic structure, usually that of an ideal in a cyclotomic number field.

While this problem is believed to be difficult for general lattice instances, it was uncovered during the past decade that one can use log-lattices from unit and S-unit groups together with new quantum algorithms to recover mildly short vectors in structured lattices. The vectors obtained by these so-called S-unit attacks are not short enough to have cryptographic consequences, but algorithms are rapidly developing and it has been claimed that an S-unit attack could recover arbitrarily short vectors.

In this work we present the mathematical background for understanding structured lattices in cryptography, we describe and survey the evolution of S-unit attacks and their relation to post-quantum cryptography. In the case of prime-power cyclotomics, we compare the behaviour of the log-unit lattice to that of a random lattice, justifying why S-unit attacks could reach much shorter vectors than currently expected. Finally, we state a recent conjecture by Daniel J. Bernstein regarding S-unit attacks.

Contents

Chapter 1

Introduction

1.1 Post-Quantum Cryptography

Classically, *cryptography* is the art of enabling secure communication in the presence of malicious eavesdroppers. Modern cryptography as a science relies on mathematical theory and computational hardness assumptions. *Public-key cryptography* enables parties to communicate securely without the need to share a secret key beforehand. Classical public-key cryptography is used daily by browsers for web authentication, by banks for secure transactions, and by instant messaging services for secure communication. Such classical systems rely on the hardness of discrete logarithm problems or integer factoring.

In 1994, Peter Shor introduces a quantum algorithm that breaks integer factoring and discrete logarithms in polynomial time, rendering all classical public-key cryptography useless against a quantum computer. This sparks the need for a new way to go about cryptography: *Post-quantum cryptography* is the area of cryptography in which security is studied under the assumption that adversaries have access to a fully-functioning quantum computer. Importantly, the user only has access to a classical computer, so post-quantum cryptography is not to be confused with quantum cryptography. Over the past couple decades, quantum computers have grown from mere intellectual curiosity to realistic not-too-distant future technology. With massive investments in the field, the prospect of a crypto-breaking scale quantum computer is now a serious threat to security. Entities have already started to intercept and store scary amounts of encrypted communications, in the hope to one day decipher them when the appropriate technology becomes available.

In 2017, the US's National Institute of Standards and Technology (NIST) releases a call for post-quantum public-key encryption and signature algorithms, hoping to find the best algorithms and standardise them. The number of candidates has gone

down from 69 in Round 1 to 15 in Round 3 in 2020, of which four have been selected for standardisation in July 2022, and four others made it to Round 4 for further discussion. This field is still at a very early stage, with new attacks being discovered regularly. This was illustrated again at the end of July 2022, when a spectacular attack was released, breaking one of the remaining Round 4 contenders [CD22].

The main propositions for post-quantum cryptography fall in one of the following categories, of which we give a very brief description:

- Code-based cryptography: uses hard problems related to error-correcting codes. Correcting errors can be easy for well-chosen codes, but is a difficult problem for random codes.

- Isogeny-based cryptography: uses hard problems related to isogenies, a special kind of morphism between elliptic curves. Explicitly computing an isogeny between elliptic curves defined over finite fields is a difficult problem.

- Hash-based cryptography: uses hash functions, that map strings of arbitrary size to strings of fixed size. Finding preimages of a given value is difficult.

- Lattice-based cryptography: uses hard problems related to lattices either for construction or in proofs of security. Finding short vectors in high-dimensional lattices is difficult.

- Multivariate cryptography: uses systems of multivariate quadratic equations. Solving these systems without a trapdoor is difficult.

This dissertation focuses on the fourth type: lattice-based cryptography. Three out of the four algorithms NIST decided to standardise are based on lattices, and they are widely considered to be the most promising for resisting against quantum computers. As a disclaimer however, deeper study is of course needed, and it would be unwise to rely only on one of the above categories.

1.2 This Work

Introduction The most efficient lattice-based schemes in cryptography rely on lattices with additional structure such as *ideal lattices*, typically lattices corresponding to ideals in families of rings, say for example $\mathbb{Z}[X]/(P(X))$ where $P(X) = X^{2^k} + 1$. One

of the main examples is the Ring-Learning With Errors (Ring-LWE) scheme, introduced in [SSTX09] and [LPR10]. [LPR10] also proves that the security of Ring-LWE relies on the worst-case hardness of Ideal-SVP: finding short vectors in ideal lattices.

Until 2014 and [CGS14] (see the next section, Section 1.3), no good attack using that extra structure was known against Ideal-SVP-like problems, and the best attacks were the same as those known for general lattices, only exploiting their additive structure. [CGS14] and [Ber14] sketch the idea behind new multiplicative attacks: *unit* and *S-unit attacks*. These attacks are quantum and exploit breakthrough results from [EHKS14] and [BS16]. Since then, a series of papers came out analysing and upgrading first unit attacks and then S-unit attacks: see [CDPR16], [CDW17], [DPW19], [PMHS19], [BRL20], [CDW21] and [BLNRL21]. It is common when analysing a lattice to rely on the heuristic assumption that the lattice behaves like what would be expected from a randomly generated instance.

Recently, Bernstein gave a talk at SIAM21 [Ber21] conjecturing that S-unit attacks can reach much better complexity than stated in the above string of works. The talk was later reproduced by Lange at ANTS-XV [Lan22], endorsing the conjecture. In [BL21] they argue that usual lattice heuristics mentioned above should not be used for analysing S-unit attacks, and state that this gives a first insight into why the conjecture is not too far-fetched. However, more evidence would be needed to draw any meaningful conclusions about the true power of S-unit attacks.

Contributions The primary objective of this dissertation is to give the reader a good understanding of unit and S-unit attacks, and their place in recent debates in cryptography. Such a survey has never been done before, as it covers fairly recent research. Following and generalising ideas from the recent paper by Bernstein and Lange [BL21], we argue that lattices used in the unit and S-unit attacks can not be analysed using heuristics that are true for random lattices. This leads us to give a first formal statement for Bernstein's conjecture.

Roadmap Chapter 1 is an introduction, including an overview of the SOLILOQUY scheme in Section 1.3. Chapters 2 and 3 cover the background material on lattices and number theory respectively, including a presentation of the quantum algorithms used in the attacks in Section 3.5. In Chapter 4 we discuss the use of structured lattice in cryptography, and justify why Ideal-SVP is so important. In Chapter 5, we give a brief overview of existing attacks on Ideal-SVP, and in Chapter 6 we describe the unit and S-unit attacks as seen in recent literature, involving reduction in the

log-unit and log-S-unit lattices. Finally, in Chapter 7, we prove that the log-unit and log-S-unit lattices do not behave like typical random lattices, and finally we state Bernstein's conjecture.

1.3 Soliloquy

To provide the reader with a concrete insight into what type of problems this dissertation tackles, and for historical reasons, we give a detailed presentation of the SOLILOQUY cryptosystem. SOLILOQUY was introduced secretly in 2007 by researchers from CESG (GCHQ's old information security arm) as a supposedly quantum-resistant key-encapsulation mechanism. A key-encapsulation mechanism or KEM is a cryptographic protocol that enables two parties to securely agree on a secret key, that can later be used for further secure communication. This KEM relies on supposed quantum hardness of some structured lattice problem, for which no attack better than those already known for general lattice problems were known. In 2014, the same team of researchers simultaneously make SOLILOQUY public and sketched a full key-recovery attack against it in [CGS14], involving a quantum algorithm. This paper is famously entitled "SOLILOQUY: A Cautionary Tale", and is the first known example of an attack on structured lattices that really uses the extra algebraic structure. It follows from breakthrough paper [EHKS14] describing new quantum algorithms for algebraic number theory, and the attack is very quickly made fully rigorous in [CDPR16] and [BS16]. Here is how it goes (the less experienced reader may want to consider reading Chapters 2 and 3 first).

Background maths: K denotes the cyclotomic field $\mathbb{Q}(\zeta_n)$ where n is a (small) prime number of size about 10 bits, and $\zeta_n = \exp\left(\frac{2i\pi}{n}\right)$ is a primitive n-th root of unity. Its ring of integers is $\mathcal{O}_K = \mathbb{Z}[\zeta_n]$. For p a (larger) prime number such that $p \equiv 1 \bmod n$, the principal ideal $(p) = p\mathcal{O}_K$ decomposes into a product of prime ideals of norm p:

$$p\mathcal{O}_K = \prod_{i=1}^{n-1} \mathfrak{p}_i. \tag{1.1}$$

The $\mathfrak{p}_i$'s are permuted by the Galois group $\mathrm{Gal}(K/\mathbb{Q}) \cong (\mathbb{Z}/n\mathbb{Z})^\times$. This and more can be seen as a consequence of a theorem of Dedekind, exposed in [ST02], Theorem 10.1. To get the decomposition of (p) we thus need to look at the factors of the minimal polynomial $\Phi_n = \frac{X^n-1}{X-1}$ of ζ_n modulo p. Because n divides $p-1$, the values $b^{\frac{p-1}{n}}$ for $b \in (\mathbb{Z}/p\mathbb{Z})^\times$ give n n-th roots of unity modulo p, are distinct and the splitting

field for Φ_n is just $\mathbb{F}_p$. Now Equation 1.1 follows and we also get that each of the $n-1$ non-trivial roots of unity mod p can be associated with one of the $\mathfrak{p}_i$. More precisely, $\mathfrak{p}_i = (p, \zeta_n - c_i)$ for c_i such a n-th root. Heuristically for a random choice of p, we expect $c = 2^{\frac{p-1}{n}}$ to be non-trivial with probability approximately $(1 - \frac{1}{n}))$, which means we can usually focus on the following factor $\mathfrak{p} = (p, \zeta_n - c)$.

Key Generation: From a fixed n, Alice generates a candidate key element

$$\alpha = \sum_{i=1}^{n} a_i \zeta_n^i \in \mathcal{O}_K,$$

where the coefficients a_i are sampled according to a discrete Gaussian distribution of mean 0 and width σ. The ideal (α) is in fact a lattice with lots of structure, as it can be represented by the following short basis.

$$\begin{pmatrix} a_1 & a_2 & \dots & a_{n-1} & a_n \\ a_n & a_1 & \dots & a_{n-2} & a_{n-1} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ a_3 & a_4 & \dots & a_1 & a_2 \\ a_2 & a_3 & \dots & a_n & a_1 \end{pmatrix}$$

To continue with the key generation procedure, Alice computes

$$p = \mathcal{N}(\alpha) = \prod_{\sigma \in \mathrm{Gal}(K/\mathbb{Q})} \sigma(\alpha),$$

and ensures that p is prime and $c = 2^{\frac{p-1}{n}}$ is not the trivial root of unity mod p, else she generates a new α. Note that from the definition of p, $p \equiv 1 \bmod n$ follows and the previous paragraph can be used freely, with the same notations. Our condition on c ensures that (α) is one of the factors $\mathfrak{p}_i$, and by taking the right Galois conjugate (there are not too many to choose from anyways as n is not big), Alice can force $(\alpha) = \mathfrak{p}$. She keeps α as her secret key, and releases p as her public key. The public key p contains all the information to recover the ideal $\mathfrak{p}$, effectively acting as a *bad basis* of the ideal lattice generated by α.

Encryption: The ideal $\mathfrak{p}$ has norm p so $\mathcal{O}_K/\mathfrak{p} \cong \mathbb{F}_p$. This induces a natural homomorphism

$$\begin{array}{rccc} \psi & : & \mathcal{O}_K & \to & \mathbb{F}_p \\ & & \sum_{i=1}^{n} e_i \zeta_n^i & \mapsto & \sum_{i=1}^{n} e_i c^i \bmod p \end{array}.$$

The objective of the KEM is for Alice and Bob to securely share a key. The function ψ will be used for encapsulating a randomly chosen key from $\mathcal{O}_K$. Bob generates a key

$$\varepsilon = \sum_{i=1}^{n} e_i \zeta_n^i \in \mathcal{O}_K$$

by sampling the coordinates e_i from a discrete Gaussian distribution with mean 0 and variance σ'. Because of how it was generated, ε is a short element of $\mathcal{O}_K$. Bob then sends $z = \psi(\varepsilon) \in \mathbb{F}_p$ to Alice.

Decryption: Suppose as in the Encryption paragraph that Bob has generated a key $\varepsilon = \sum_{i=1}^{n} e_i \zeta_n^i \in \mathcal{O}_K$, and has sent $z = \psi(\varepsilon)$ to Alice. She would like to recover ε from its image $z \in \mathbb{F}_p$. This equates to recovering the right coset in the quotient $\mathcal{O}_K/\mathfrak{p}$. As $\varepsilon \in z + \alpha\mathcal{O}_K$ and ε is short, Alice must solve an instance of the CVP in the ideal lattice $\alpha\mathcal{O}_K$ (see sections 2.2 and 3.3). Given that she knows a short basis α for the ideal, she can use Babai's round-off algorithm to do exactly that. This simply means that

$$\varepsilon = z - \lceil z\alpha^{-1} \rfloor \cdot \alpha,$$

provided ε was initially chosen small enough, and where $\lceil \cdot \rfloor$ denotes coordinate-wise rounding to the nearest integer. For a proper proof of correctness, see Smart and Vercauteren's Fully-Homomorphic encryption scheme, that uses a similar procedure [SV09]. Alice and Bob now share the randomly generated key ε, and can use it for other cryptographic purposes.

Discussion: We shall see later why SOLILOQUY is completely broken. In a few words, note that p, n, c are public, meaning that an attacker can construct the ideal $(\alpha) = (p, \zeta_n - c)$. Seeing (α) as a lattice, the scheme relies on the fact that it should be hard for an attacker to find a short basis of this lattice that would enable the rounding procedure. Here, a complete break would be obtained if an attacker could recover a shortest generator of the principal ideal (α). Therefore, the supposedly hard problem is: given a basis of an ideal *guaranteed to be principal and have a short generator*, find any short generator of the ideal. This problem is called the Short-Generator Principal Ideal Problem (SG-PIP for short). The guarantee in italic is in fact what makes SOLILOQUY horribly insecure.

Chapter 2

Lattice Background

2.1 Basic Facts

Definition 2.1.1. A lattice of $\mathcal{L} \subset \mathbb{R}^n$ is a discrete additive subgroup of $\mathbb{R}^n$.

The word *discrete* is used here in the context of the usual Euclidean topology, that is for any point $v \in L$, there exists a small distance $\varepsilon > 0$ such that no $u \in \mathcal{L}$ different to v satisfies $\|u - v\| \leq \varepsilon$.

Example 2.1.2. $\mathcal{L} = \mathbb{Z}^k$ for $k \leq n$ is a lattice of $\mathbb{R}^n$.

Example 2.1.3. The lattices of $\mathbb{R}$ are exactly the $\mathbb{Z}$-vector spaces $\alpha\mathbb{Z}$ for $\alpha \in \mathbb{R}_{\geq 0}$.

It is important in computer science to know how to represent a lattice. Even though lattices usually have countably many elements, it is sufficient to know a $\mathbb{Z}$-basis.

Proposition 2.1.4. *Let $\mathcal{L}$ be a lattice of $\mathbb{R}^n$, then there exists an $m \leq n$ and m independent vectors $b_1, \ldots, b_m \in \mathbb{R}^n$ such that*

$$\mathcal{L} = \left\{ \sum_{i=1}^{m} x_i b_i \;\middle|\; x_1, \ldots, x_m \in \mathbb{Z} \right\}.$$

In this case we say that $B = (b_1, \ldots, b_m)$ is a basis of $\mathcal{L}$ and we use the notations $\mathcal{L}(B)$ or $\mathcal{L}(b_1, \ldots, b_m)$ to denote the lattice. Moreover, m is called the dimension of $\mathcal{L}$, and if $m = n$, we say that $\mathcal{L}$ is full-rank.

Therefore, any full-rank lattice can be represented by the matrix $B \in GL_n(\mathbb{R})$ whose columns are $(b_1, \ldots, b_n)$. Reciprocally, given an invertible matrix $B \in GL_n(\mathbb{R})$, the set $\{Bx \mid x \in \mathbb{Z}^n\}$ defines a full-rank lattice. However, like a vector space has multiple bases, a lattice has multiple representations:

Theorem 2.1.5. *Let $B_1, B_2 \in GL_n(\mathbb{R})$ represent the lattices $\mathcal{L}(B_1), \mathcal{L}(B_2)$, then $\mathcal{L}(B_1) = \mathcal{L}(B_2)$ if and only if there exists a matrix $U \in \mathbb{Z}^{n\times n}$ such that $\det(U) = \pm 1$ and $B_1 = B_2 U$.*

Proof. First suppose $\mathcal{L}(B_1) = \mathcal{L}(B_2)$, then columns of B_1 are in $\mathcal{L}(B_2)$. Therefore there exists a matrix $U_1 \in \mathbb{Z}^{n\times n}$ such that $B_1 = B_2 U_1$. By the same argument, there exists $U_2 \in \mathbb{Z}^{n\times n}$ such that $B_1 U_2 = B_2$. Combining both identities we get $B_1 U_2 U_1 = B_2 U_1 = B_1$, from which $U_2 U_1 = I_n$. However U_1 and U_2 are integeral matrices and thus have integer determinant, but $\det(U_2)\det(U_1) = 1$ so $\det(U_1) = \pm 1$. Second suppose $B_1 = B_2 U$ with $U \in \mathbb{Z}^{n\times n}$ and $\det(U) = \pm 1$, then U is invertible and has inverse $U^{-1} = \det(U)^{-1}\operatorname{Co}(U)^T$ (where $\operatorname{Co}(U)$ is the matrix of cofactors of U). $\operatorname{Co}(U)$ is integral and $\det(U) = \pm 1$ so U^{-1} also has integer coefficients. Now since U is integral, vectors spanned by B_1 are spanned by B_2, so $\mathcal{L}(B_1) \subset \mathcal{L}(B_2)$, and since U^{-1} is integral, $\mathcal{L}(B_2) \subset \mathcal{L}(B_1)$, which concludes. □

Corollary 2.1.6. *The set of full-rank lattices of $\mathbb{R}^n$ is $GL_n(\mathbb{R})/GL_n(\mathbb{Z})$.*

Definition 2.1.7. The *volume* of a lattice $\mathcal{L}$ is the absolute value of the determinant of any basis $(b_1, \ldots, b_n)$ of $\mathcal{L}$:

$$\operatorname{vol}(\mathcal{L}) = |\det(b_1, \ldots, b_n)|.$$

This definition is justified by Theorem 2.1.5, as if B_1 and B_2 are two bases for $\mathcal{L}$, then there exists $U \in \mathbb{Z}^{n\times n}$ such that $|\det(U)| = 1$ and $B_1 = B_2 U$, whence $|\det(B_1)| = |\det(B_2)|$.

2.2 Finding Short Vectors

We have seen that all lattices have multiple possible bases. While they all consist of the same number of elements, some of them are better for different reasons. Picture the lattice $\mathbb{Z}^n$, it is very easy to decompose any $x \in \mathbb{Z}^n$ as a linear combination of the canonical basis vectors, however this problem looks much harder if we have a basis consisting of only vectors that roughly point in the same direction. This justifies why from the point of view of a computer scientist (or a cryptographer), one would prefer knowing a basis consisting of only orthogonal vectors. In dimensions two and above, an orthogonal basis does not usually exist, so instead, we look for very small vectors. Geometrically, the smallest vectors are usually close to orthogonal. This is the reason why ultimately, our goal is to look for short vectors, so that we can understand the lattice with as little computing time as possible.

Definition 2.2.1. For the Euclidean norm, the shortest non-zero vector of a lattice $\mathcal{L}$ of dimension $n \geq 1$ exists and its norm is denoted $\lambda_1(\mathcal{L})$.

Proof. $\dim(\mathcal{L}) \geq 1$ so there exists a non-zero vector $u \in \mathcal{L}$. The set defined by $\{v \in \mathcal{L} \mid 0 < \|v\| \leq \|u\|\}$ is compact, discrete and non-empty, therefore it is finite and has a minimum, $\lambda_1(\mathcal{L})$. Note that the vector that reaches the minimum is not unique, for exemple, if u is a shortest vector, then so is $-u$. □

This leads us to consider the following algorithmic problems, relatively to the Euclidean norm (see [Ngu10] for a more precise overview):

Problem 2.2.2 (Shortest Vector Problem (SVP))**.** *Given a basis of a lattice $\mathcal{L}$, find a nonzero vector $u \in \mathcal{L}$ such that $\|u\| = \lambda_1(\mathcal{L})$.*

Problem 2.2.3 (Approximate Shortest Vector Problem (γ-SVP))**.** *Given a basis of a lattice $\mathcal{L}$ and an approximation factor $\gamma \geq 1$, find a non-zero vector $u \in \mathcal{L}$ such that $\|u\| \leq \gamma\lambda_1(\mathcal{L})$.*

Problem 2.2.4 (Closest Vector Problem (CVP))**.** *Given a basis of a full-rank lattice $\mathcal{L} \subset \mathbb{R}^n$ and a point $v \in \mathbb{R}^n$, find a vector $u \in \mathcal{L}$ such that $\|u - v\| = \operatorname{dist}(v, \mathcal{L})$.*

Proof of existence of the closest vector can be adapted from the proof of existence of the shortest vector.

Problem 2.2.5 (Approximate Closest Vector Problem (γ-CVP))**.** *Given a basis of a full-rank lattice $\mathcal{L} \subset \mathbb{R}^n$, a point $v \in \mathbb{R}^n$ and an approximation factor $\gamma \geq 1$, find a vector $u \in \mathcal{L}$ such that $\|u - v\| \leq \gamma \operatorname{dist}(v, \mathcal{L})$.*

SVP was first proved to to be NP-hard for randomised reductions [Ajt98], and later [Mic01] proved that γ-SVP for any $\gamma < \sqrt{2}$ is NP-hard. Clearly, CVP is a generalisation of SVP, and the approximation problems 1-SVP and 1-CVP are exactly the search problems SVP and CVP respectively. Moreover, [GMSS99] proves that any hardness for SVP implies the same for CVP. This makes SVP and CVP very natural choices for the design of cryptosystems.

A practical method for finding short vectors in lattices is the Lenstra-Lenstra-Lovász algorithm (LLL). LLL solves γ-SVP in polynomial time for an exponential (in the dimension of the lattice) approximation factor. In what follows we present the algorithm and its correctness. For the analysis of the runtime and insight on the importance of this algorithm across mathematics and computer science, see the original paper [LLL82] and the survey book [NV09]. We first recall the Gram-Schmidt orthogonalisation process, where $\langle \cdot, \cdot \rangle$ denotes the Euclidean scalar product.

Definition 2.2.6. The Gram-Schmidt orthogonalisation of n linearely independant vectors $(b_1, \ldots, b_n)$ is defined by $b_i^\star = b_i - \sum_{j=1}^{i-1} \mu_{i,j} b_j^\star$, where $\mu_{i,j} = \frac{\langle b_i, b_j^\star \rangle}{\|b_j^\star\|^2}$.

This process generates a new orthogonal basis by successively projecting each vector orthogonally to the space spanned by the previously obtained vectors.

Definition 2.2.7. Let $\frac{1}{4} < \delta < 1$, a basis $B = (b_1, \ldots, b_n) \in \mathbb{R}^n$ is δ-reduced if the following points are true:

1. $\forall 1 \leq i \leq n, j < i, |\mu_{i,j}| \leq \frac{1}{2}$;
2. $\forall 1 \leq i \leq n, \delta \|b_i^\star\|^2 \leq \|\mu_{i+1,i} b_i^\star + b_{i+1}^\star\|^2$

To understand this definition, we can look at B in the Gram-Schmidt basis and we get

$$\begin{pmatrix} \|b_1^\star\| & \mu_{1,2}\|b_1^\star\| & \ldots & \mu_{1,n}\|b_1^\star\| \\ 0 & \|b_2^\star\| & & \vdots \\ \vdots & & \ddots & \\ 0 & \ldots & & \|b_n^\star\| \end{pmatrix},$$

where the first condition says that the off-diagonal coefficients are less than a half of the corresponding diagonal coefficient on the row; and the second that when looking at the 2×2 submatrices on the diagonal, the norm of their second column is δ-close to the norm of their first.

Proposition 2.2.8. *Let $B = (b_1, \ldots, b_n)$ be a δ-reduced basis for a $\frac{1}{4} < \delta < 1$, then*

$$\|b_1\| \leq \left(\frac{2}{\sqrt{4\delta - 1}}\right)^{n-1} \lambda_1(\mathcal{L}(B)).$$

If $\delta = \frac{3}{4}$, this gives $\|b_1\| \leq 2^{\frac{n-1}{2}} \lambda_1(\mathcal{L}(B))$.

Proof. First we get from condition (2) in the definition of a δ-reduced basis that $\forall 1 \leq i \leq n, (\delta - \mu_{i+1,i}^2)\|b_i^\star\|^2 \leq \|b_{i+1}^\star\|^2$. Combining with condition (1) and iterating:

$$\|b_n^\star\|^2 \geq \left(\delta - \frac{1}{4}\right) \|b_{n-1}^\star\|^2 \geq \ldots \geq \left(\delta - \frac{1}{4}\right)^{n-1} \|b_1^\star\|^2 = \left(\delta - \frac{1}{4}\right)^{n-1} \|b_1\|^2,$$

from which

$$\|b_1\| \leq \left(\delta - \frac{1}{4}\right)^{\frac{n-1}{2}} \min_i \|b_i^\star\|.$$

We finally prove that $\lambda_1(\mathcal{L}(B)) \geq \min_i \|b_i^\star\|$, which concludes: let $u \in \mathcal{L}(B)$ be a nonzero lattice vector, then there exists a $1 \leq k \leq n$ and integers $(a_i)_{1 \leq i \leq k}$ with $a_k \neq 0$ such that $u = \sum_{i=1}^{k} a_i b_i$. So

$$u = \sum_{i=1}^{k} a_i \left(b_i^\star + \sum_{j=2}^{i} \mu_{i,j} b_j^\star \right) = a_k b_k^\star + S,$$

with $\langle b_k^\star, S \rangle = 0$. Hence $\|u\| = \|a_k b_k^\star\| + \|S\| \geq \|b_k^\star\|$. From this we have our conclusion. □

This proposition means that from a δ-reduced basis, we get an exponential solution to γ-SVP. We now present LLL itself in Algorithm 1.

Algorithm 1: The Lenstra-Lenstra-Lovász algorithm, with parameter $\delta \in (1/4, 1)$

input : Basis $(b_1, \ldots, b_n) \in \mathbb{R}^n$
output: A δ-reduced basis of $\mathcal{L}(B)$

Gram-Schmidt process;
Compute $b_1^\star, \ldots, b_n^\star$;
Reduction phase;
for $i \leftarrow 2$ **to** n **do**
  for $j \leftarrow i - 1$ **to** 1 **do**
    $b_i \leftarrow b_i - \lceil \langle b_i, b_j^\star \rangle / \|b_j^\star\|^2 \rfloor b_j$;
Swap phase;
if $\exists i \quad s.t. \quad \delta \|b_i^\star\|^2 \leq \|\mu_{i+1,i} b_i^\star + b_{i+1}^\star\|^2$ **then**
  $b_i \leftrightarrow b_{i+1}$;
  Go back to Reduction phase;
Output $(b_1, \ldots, b_n)$;

Here $\lceil \cdot \rfloor$ denotes rounding to the nearest integer.

Lemma 2.2.9. *If LLL terminates, it produces a δ-reduced basis.*

Proof. Condition (2) is clearly satisfied as it is directly taken care of by the swap phase. The basis remains a basis as it is only modified through transvections and swaps. Its Gram-Schmidt orthogonalisation also stays the same throughout. Remains to prove that the reduction phase ends with all $|\mu_{i,j}| \leq \frac{1}{2}$. Let $i > j$ and look at the step i of the first loop, step j of the second. The algorithm is designed to subtract column j to column i enough times to small off-diagonal terms, in fact we have

$$|\mu_{i,j}| = \|b_j^\star\|^{-2} \left| \left\langle b_i - \left\lceil \frac{\langle b_i, b_j^\star \rangle}{\|b_j^\star\|^2} \right\rfloor b_j, b_j^\star \right\rangle \right| = \left| \frac{\langle b_i, b_j^\star \rangle}{\|b_j^\star\|^2} - \left\lceil \frac{\langle b_i, b_j^\star \rangle}{\|b_j^\star\|^2} \right\rfloor \frac{\langle b_j, b_j^\star \rangle}{\|b_j^\star\|^2} \right| \leq \frac{1}{2},$$

where we used the projection identity $\langle b_j, b_j^\star \rangle = \langle b_j^\star, b_j^\star \rangle = \|b_j^\star\|^2$. Therefore, after step (i, j) condition (1) is and stays valid, and the concludes our proof. □

We will later see in section 5.1 how LLL can be generalised to allow for a better approximation factor in γ-SVP.

An approach for solving γ-CVP is the following.

Definition 2.2.10 (Dual basis)**.** Let $B = (b_1, \ldots, b_n)$ be a basis of $\mathbb{R}^n$. The basis $B^\vee = (b_1^\vee, \ldots, b_n^\vee)$ where $\langle b_i^\vee, b_j \rangle = \delta_{ij}$ is called the dual basis of B. δ_{ij} is the Kronecker symbol of i, j.

Definition 2.2.11 (Babai's Round-off algorithm)**.** Let $\mathcal{L}(B)$ be a lattice of $\mathbb{R}^n$, and $v \in \mathbb{R}^n$ a target vector, return $B \cdot \lceil (B^\vee)^T \cdot v \rfloor$, where here $\lceil \cdot \rfloor$ denotes rounding of each coordinate to the nearest integer.

The following fact is standard and tells us when the round-off algorithm works.

Proposition 2.2.12. *Let $B = (b_1, \ldots, b_n)$ be a basis of $\mathbb{R}^n$, and let $v = u + e \in \mathbb{R}^n$ for some $u \in \mathcal{L}(B)$ and $e \in \mathbb{R}^n$. Suppose $\frac{1}{2} \leq \langle b_j^\vee, e \rangle < \frac{1}{2}$ for all $j \in \{1, \ldots, n\}$. Then Babai's round-off algorithm with input (B, v) outputs u.*

Proof. $u \in \mathcal{L}(B)$ so there exists a $z \in \mathbb{Z}^n$ such that $u = Bz$. The product $(B^\vee)^T B$ is the identity matrix so multiplying $v = u + e$ by $(B^\vee)^T$ gives $(B^\vee)^T v = z + (B^\vee)^T e$. Moreover, the coordinates of $(B^\vee)^T e$ are $\langle b_j^\vee, e \rangle$, therefore the round-off algorithm outputs $B \cdot \lceil (B^\vee)^T v \rfloor = Bz = u$. □

Another slightly more expensive way to solve γ-CVP is Babai's nearest plane algorithm, presented in Algorithm 2.

Algorithm 2: Babai's nearest plane algorithm

input : Basis $B = (b_1, \ldots, b_n) \in \mathbb{R}^n$, target vector $v \in \mathbb{R}^n$
output: A vector $u \in \mathcal{L}(B)$ close to v.

Compute $b_1^\star, \ldots, b_n^\star$;
Run LLL on B with $\delta = \frac{3}{4}$;
$w \leftarrow v$;
for $i \leftarrow n$ **to** 1 **do**
  $w \leftarrow w - \lceil \langle w, b_i^\star \rangle / \langle b_i^\star, b_i^\star \rangle \rfloor b_i$;
Output $u = v - w$;

Lemma 2.2.13. *Algorithm 2 with input lattice $\mathcal{L}(B)$ and target $v \in \mathbb{R}^n$ terminates in polynomial time and outputs a vector $u \in \mathcal{L}$ that satisfies*

$$\|u - v\|^2 \leq \frac{1}{4} \sum_{i=1}^{n} \|b_i^\star\|^2.$$

Proof. The runtime is just the runtime of LLL, followed by a linear loop. Therefore the total runtime is polynomial. Now each step i of the loop ensures that the i-th coordinate of w is smaller than $\frac{1}{2}\|b_i^\star\|^2$ by the rounding process. Therefore

$$\|u - v\|^2 = \|w\|^2 \leq \frac{1}{4} \sum_{i=1}^{n} \|b_i^\star\|^2.$$

□

With a bit more work we can prove that Babai's nearest plane algorithm solves $2^{\frac{n}{2}}$-CVP. See [Bab86] for the original analysis.

2.3 Random Lattices

In order to work with lattices in practice it would be nice to have a way to generate random instances. However it is not an easy task to randomly sample integers. In this subsection, we prove the remarkable fact that there exists a canonical way to choose lattices of $\mathbb{R}^n$ with fixed determinant at random. This is a result from Siegel [Sie45], whose proof can be skipped on a first read.

Theorem 2.3.1. *There exists a unique probability measure μ on the space*

$$\mathrm{SL}_n(\mathbb{R})/\,\mathrm{SL}_n(\mathbb{Z}),$$

invariant by $\mathrm{SL}_n(\mathbb{R})$, ie such that for all $M \in \mathrm{SL}_n(\mathbb{R})$ and $S \subset \mathrm{SL}_n(\mathbb{R})/\,\mathrm{SL}_n(\mathbb{Z})$, $\mu(MS) = \mu(S)$.

Proof. $\mathrm{SL}_n(\mathbb{R})$ is a locally compact Hausdorff topological group, and therefore there exists a unique (up to a positive multiplicative constant) Haar measure. The measure we are looking for on determinant-1 lattices is essentially a projection of the Haar integral on $\mathrm{SL}_n(\mathbb{R})$ with the correct normalisation factor. In order to prove it passes well to the quotient and is unique there, we need to introduce a bit of theory on topological groups. If G is a locally compact Hausdorff group, μ a left-invariant Haar measure on G and $g \in G$, then the map $E \mapsto \mu(Eg)$ defined on the Borel sets on

G also defines a left-invariant Haar measure on G. By unicity up to a scalar, there exists a constant $\Delta_G(g)$ such that

$$\mu(Eg) = \Delta_G(g)\mu(E)$$

for all E. This defines the group homomorphism $\Delta_G : g \in G \mapsto \Delta_G(g) \in \mathbb{R}_{>0}$, called the modular function of G. If Δ_G is identically 1, G is said to be unimodular. In particular, discrete groups are always unimodular. We need the following Theorem:

Theorem 2.3.2. *Let G be a locally compact Hausdorff group, and H a closed subgroup of G. Then a left G-invariant Haar measure on G/H is unique up to positive scalar multiples, and exists if and only if*

$$\Delta_G(h) = \Delta_H(h)$$

for all $h \in H$.

Proof. See [AM07], Theorem 2.3.5. □

$SL_n(\mathbb{Z})$ is a discrete closed subgroup of $SL_n(\mathbb{R})$, therefore it is unimodular. We will now show that (1) $SL_n(\mathbb{R})$ is unimodular, and that (2) There is a set F of finite Haar measure such that $F \cdot SL_n(\mathbb{Z}) = SL_n(\mathbb{R})$. From (1), using Theorem 2.3.2 we get existence and uniqueness up to a scalar multiple of the invariant measure on $SL_n(\mathbb{R})/SL_n(\mathbb{Z})$; and from (2) we get that we can normalise said measure into a unique invariant probability measure.

Proof of (1): Given A and B topological groups with left Haar measures da and db respectively, the product $da \times db$ on $A \times B$ is a left-invariant measure that has to be the Haar measure. Moreover, the modular function of $A \times B$ is the product of the modular functions of A and B. The Haar measure on $GL_n(\mathbb{R})^+$, real matrices with positive determinants is unimodular and given by $d\mu(g) = \frac{dg}{(\det g)^n}$, so the Haar measure on $\mathbb{R}_{>0}$ is $d\mu(g) = \frac{dg}{g}$ and unimodular by taking $n = 1$. By the topological group isomorphism $GL_n(R)^+ \to SL_n(R) \times \mathbb{R}_{>0}$, we deduce that $SL_n(R)$ is in fact unimodular.

Proof of (2): We first prove that $SL_n(\mathbb{R}) = \mathcal{S}_{\frac{2}{\sqrt{3}},\frac{1}{2}} \cdot SL_n(\mathbb{Z})$ where $\mathcal{S}_{\frac{2}{\sqrt{3}},\frac{1}{2}}$ refers to a Siegel set, defined as follows:

Definition 2.3.3 (Siegel set)**.** Let $A > 0$ and $B > 0$, the Siegel set $\mathcal{S}_{A,B}$ is the set of elements $g \in SL_n(\mathbb{R})$ such that if $g = kau$ is the Iwasawa decomposition of g,

$$\forall 1 \leq i < j \leq n : \frac{a_i}{a_{i+1}} \leq A \text{ and } |u_{ij}| \leq B,$$

where the Iwasawa decomposition refers to the following:

Proposition 2.3.4 (Iwasawa decomposition)**.** *Let $g \in SL_n(\mathbb{R})$, then there exists a unique decomposition*

$$g = kau,$$

where $k \in SO_n(\mathbb{R})$, a is a diagonal matrix with determinant 1 and positive entries, and u is upper triangular with only 1's on the diagonal.

Proof. From $(v_1, \ldots, v_n)$ the columns of g, derive $\tilde{u} = (\tilde{v}_1, \ldots, \tilde{v}_n)$ and $\tilde{a}$ diagonal from the Gram-Schmidt orthonormalisation process such that $g\tilde{u}\tilde{a}$ is orthonormal, where $\tilde{u}$ is upper triangular with 1's on the diagonal, and $\tilde{a}$ must have determinant 1. Define $k = g\tilde{u}\tilde{a}$, $u = \tilde{u}^{-1}$ and $a = \tilde{a}^{-1}$ to conclude. □

We will need the following lemma:

Lemma 2.3.5. *Let $L \subset \mathbb{R}^n$ be a lattice, then there exists a basis $v_1, \ldots, v_n$ of L such that*

$$\forall 2 \leq i \leq n, \|v_i'\|^2 \geq \frac{3}{4}\|v_{i-1}^2\|,$$

where v_i' denotes the orthogonal projection of v_i on $\langle v_1, \ldots, v_{i-1}\rangle^{\perp}$.

Proof. Apply LLL with parameter $\delta = \frac{3}{4}$. □

We can now prove that any $g \in SL_n(\mathbb{R})$ can be pushed into the fundamental domain $\mathcal{S}_{\frac{2}{\sqrt{3}},\frac{1}{2}}$ by multiplication by an element of $SL_n(\mathbb{Z})$. Fix a $g \in SL_n(\mathbb{R})$, it corresponds to a lattice of $\mathbb{R}^n$ with basis the columns of g. By Lemma 2.3.5 there exists a basis $(v_1, \ldots, v_n)$ of L such that for $2 \leq i \leq n$ we have

$$\|v_i'\|^2 \geq \frac{3}{4}\|v_{i-1}^2\|.$$

Therefore there is a $\alpha \in SL_n(\mathbb{Z})$ such that $g\alpha$ corresponds to this new basis. Looking at its Iwasawa decomposition $g\alpha = kav$ with for $2 \leq i \leq n$,

$$\frac{a_{i-1}}{a_i} = \frac{\|v_{i-1}'\|}{\|v_i'\|} \leq \frac{\|v_{i-1}\|}{\|v_i'\|} \leq \sqrt{\frac{4}{3}} = \frac{2}{\sqrt{3}}.$$

We can then choose a $\beta \in SL_n(\mathbb{Z})$ in a way that all off-diagonal coefficients of $u = v\beta$ satisfy $|u_{ij}| \leq \frac{1}{2}$ for all $i < j$. Thus

$$g(\alpha\beta) \in \mathcal{S}_{\frac{2}{\sqrt{3}},\frac{1}{2}},$$

proving the first part of (2). In order to complete the proof of (2), we show that Siegel sets have finite measure for the Haar measure on $SL_n(\mathbb{R})$. For $g = kau$ in $SL_n(\mathbb{R})$,

using unimodularity and Theorem 2.3.2, [Lan85] Chapter 3 paragraph 1 proves that the Haar measure on $SL_n(\mathbb{R})$ is given by

$$dg = (dk)(du)(da)$$

where dk is the Haar measure on $SO_n(\mathbb{R})$, du on the group of upper triangular matrices with 1's on the diagonal, and da on positive diagonal matrices of determinant 1. Note that da and du are reversed compared to the usual decomposition. This accounts for an extra factor $\prod_{i<j} \frac{a_i}{a_j}$ obtained by looking at the decomposition $k(aua^{-1})a$ instead of kau. $SO_n(\mathbb{R})$ is compact so it has finite measure, so is the set of upper triangular matrices of $\mathbb{R}$ with 1's on the diagonal and coefficients in $[-B, B]$ (here we had $B = \frac{1}{2}$), therefore we must compute the measure of the center factor, which is:

$$\int_{\frac{a_i}{a_{i+1}} \leq A} \left(\prod_{i<j} \frac{a_i}{a_j} \right) \frac{da_1}{a_1} \dots \frac{da_{n-1}}{a_{n-1}}.$$

Using the change in variable $t_i = \frac{a_i}{a_{i+1}}$ for $1 \leq i < n$, and using that for $i < j$ $\frac{a_i}{a_j} = \prod_{k=i}^{j-1} \frac{a_k}{a_{k+1}}$, there exists non-negative integers m_i such that the center factor becomes

$$\int_{t_i \leq A} t_1^{m_1} \dots t_{n-1}^{m_{n-1}} dt_1 \dots dt_{n-1} = \prod_{i=1}^{n-1} \int_0^A t^{m_i} dt.$$

This is finite, and so $SL_n(\mathbb{R})/SL_n(\mathbb{Z})$ has finite measure, and this completes the proof. □

This invariant measure is very nice in theory, but less so for experimental purpose. One way to work around this difficulty is to look at generation of random integer lattices. Given integers V and n, there are finitely many n-dimensional lattices with volume V. [GM03] proves that sampling uniformly at random one of these integer lattices, and then rescaling it by $V^{1/n}$ converges asymptotically with V towards the invariant measure for volume 1 real lattices. As explained in the first section of [GM03], this sampling process is particularly simple when V is a prime number.

2.4 Gaussian Heuristic

Studying the behaviour of random lattices is interesting, as it provides heuristic guidance as to what we can assume when we encounter a lattice. Assuming our lattice $\mathcal{L}$ can be taken at random, this gives tools for estimating quantities such as the length $\lambda_1(\mathcal{L})$, the number of points in $L \cap X$ for a well-behaved set X, and this can be used for analysing how well some lattice reduction algorithms such as LLL perform. The

main assumption used time and time again in the literature and referred to as the *Gaussian heuristic* can be stated as follows:

Heuristic 2.4.1 (Gaussian heuristic)**.** *Given a lattice $\mathcal{L}$ and a "well-behaved" set X, the Gaussian Heuristic predicts the number of points in $\mathcal{L} \cap X$ to be*

$$\#(\mathcal{L} \cap X) \approx \frac{\mathrm{vol}(X)}{\mathrm{vol}(\mathcal{L})},$$

where "well-behaved" usually means Borel-measurable and convex.

It is important to note that this is just a heuristic, and in some cases is very far from the truth. It can be made more precise and/or more rigorous, but most of the time is stated as is. [BL21], Section 1.6 criticises heavily the frequent lack of clarity in the literature around defining and using the heuristic. Instead, it proposes the *Spherical model* of a lattice, as a way to model the lattice that disregards any additive structure and only keeps information about its dimension and volume.

Definition 2.4.2 (Spherical model, [BL21] Definition 3.3)**.** Let $\mathcal{L}$ be a dimension n lattice. For $j \in \mathbb{N}_{>0}$, let μ_j be a uniform random element of

$$\left\{ x \in \mathrm{span}_{\mathbb{R}}(\mathcal{L}) \mid \|x\|^n = 2j\pi^{-\frac{n}{2}}\Gamma\left(\frac{n}{2}+1\right)\mathrm{vol}(\mathcal{L}) \right\};$$

then the set $\{0\} \cup \{\pm\mu_j \mid j \in \mathbb{N}_{>0}\}$ is a spherical model of $\mathcal{L}$. Γ denotes the usual Gamma function for which $\Gamma(k) = (k-1)!$ when k is a positive integer.

In order to explain where this model comes from, we need to know the volume of a ball.

Definition 2.4.3. Let $d \in \mathbb{N}$ and $r \in \mathbb{R}_{>0}$. The d-dimensional ball of radius r is defined as

$$rB_d = \{x \in \mathbb{R}^d : \|x\|^n \leq r\}.$$

Lemma 2.4.4. *Let $d \in \mathbb{N}$ and $r \in \mathbb{R}_{>0}$. The d-dimensional volume of rB_d is*

$$\mathrm{vol}(rB_d) = \frac{r^d \pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2}+1)}.$$

Proof. This is a standard result, obtained by integration. □

From Lemma 2.4.4 and Definition 2.4.4, one can verify that the Gaussian heuristic holds for $X = rB_d$, and this justifies the *spherical* part. In a Spherical model, the length of the shortest non-zero vector μ_1 is fixed directly by the dimension and volume of the lattice. It is not unjustified to do so, because of the following result from [Rog56].

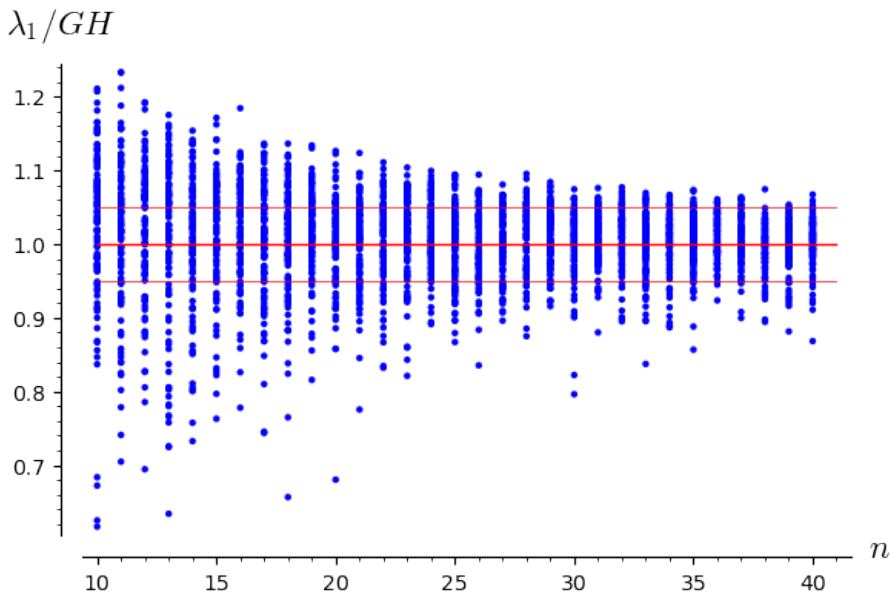

Figure 2.1: λ_1 in random lattices vs Gaussian heuristic prediction

Theorem 2.4.5 ([Rog56], Theorem 3)**.** *Let X be any Borel set in dimension n, with measure* $\mathrm{vol}(X)$*. The distribution of the number of (non-zero) pairs of points $\pm x$ of a lattice $\mathcal{L}$ in X, where $\mathcal{L}$ is sampled according to the invariant distribution has asymptotic distribution the Poisson distribution of parameter* $\frac{1}{2}\mathrm{vol}(X)$*.*

In a Gaussian Heuristic setting, one would sometimes argue that this behaviour is the expected behaviour, and therefore it can be applied to any lattice, heuristically. But why does this justify Definition 2.4.2?
The mean of a Poisson distribution is its parameter, therefore in order for the mean to be exactly one (this corresponds to the minimal non-zero vector) when a volume-one lattice $\mathcal{L}$ of dimension n intersects a ball of radius r, from Lemma 2.4.4, we would need $\frac{1}{2}\frac{r^n\pi^{n/2}}{\Gamma(n/2+1)} = 1$, so $r = (2\pi^{-n/2}\Gamma(n/2+1))^{1/n}$. Obviously this equality is not rigorous, however it can be proven with this Poisson setting that

$$\lambda_1(\mathcal{L}) = \left(1 + O\left(\frac{n}{\log(n)}\right)\right)\left(2\pi^{-\frac{n}{2}}\Gamma\left(\frac{n}{2}+1\right)\right)^{1/n}$$

with probability $1 - o(1)$ when $n \to \infty$ (see this as a consequence for balls of [SS16], Theorem 1.1). After scaling by $\mathrm{vol}(\mathcal{L})$ in the general case, we get a justification for the minimal length of a spherical model.

Lemma 2.4.6. *The length of the minimal non-zero vector in a spherical model of a lattice of volume* $\mathrm{vol}(L)$ *in dimension* n *is asymptotically equal to*

$$\sqrt{\frac{n}{2\pi e}}\,\mathrm{vol}(L)^{\frac{1}{n}}$$

when $n \to \infty$.

Proof. Given a spherical model of $\mathcal{L}$, clearly, $\min_j \|\mu_j\| = \|\mu_1\|$. From Stirling's formula for the gamma function,

$$\Gamma(n/2+1) \sim \sqrt{2\pi\frac{n}{2}}\left(\frac{n/2}{e}\right)^{\frac{n}{2}},$$

from which we deduce for $n \to \infty$

$$\begin{aligned}\|\mu_1\| &= 2^{-\frac{1}{n}}\pi^{-\frac{1}{2}}\Gamma\left(\frac{n}{2}+1\right)^{-\frac{1}{n}}\mathrm{vol}(\mathcal{L})^{\frac{1}{n}}\\ &\sim 2^{-\frac{1}{n}}\pi^{-\frac{1}{2}}(\pi n)^{-\frac{1}{2n}}\sqrt{\frac{n}{2e}}\,\mathrm{vol}(\mathcal{L})^{\frac{1}{n}}\\ &\sim \sqrt{\frac{n}{2\pi e}}\,\mathrm{vol}(L)^{\frac{1}{n}}\end{aligned}$$

□

Figure 2.1 compares the value of λ_1 predicted by the Gaussian heuristic with the actual length of the shortest vector for 100 randomly sampled lattices of dimension n for all $n \in \{10, \ldots, 40\}$. All lattices are sampled using the method described at the end of Section 2.3, with 200-bit primes. This confirms that even for fairly small dimensions, the Gaussian heuristic is accurate in predicting the size of the shortest vector. Experiments were conducted with SageMath [The22].

In Chapter 7, we will apply predictions from Gaussian heuristic/Spherical models to the log-unit and log-S-unit lattices (see Chapter 3 for definitions).

Chapter 3

Number Theory Background

Let K denote a number field of degree n over $\mathbb{Q}$, and $\mathcal{O}_K$ its ring of integers, ie the ring of algebraic integers contained in K. Given an element $\alpha \in K$, the algebraic trace and norm of α are the trace and determinant of the multiplication-by-α endomorphism $x \mapsto \alpha x$ of K, seen as a $\mathbb{Q}$-vector space. The trace is denoted $\mathrm{Tr}(\alpha)$ and the norm $\mathcal{N}(\alpha)$. For any $\mathbb{Z}$-basis $(\omega_1, \ldots, \omega_n)$ of $\mathcal{O}_K$, the value of $\det\left(\mathrm{Tr}(\omega_i\omega_j)\right)_{i,j}$ is the same and is called the discriminant of K, denoted Δ_K. Loosely speaking, $|\Delta_K|$ measures the size of the number field, and for this reason most complexities in computational algebraic number theory depend on $\ln|\Delta_K|$. Units in $\mathcal{O}_K$ are the group $\mathcal{O}_K^\times$ of elements of $\mathcal{O}_K$ that have norm 1.

3.1 Cyclotomic Number Fields

Let $\zeta_n = \exp(2i\pi/n)$, the first primitive n-th root of unity. We define the n-th cyclotomic field as the number field $\mathbb{Q}(\zeta_n)$.

Proposition 3.1.1. *Let $K = \mathbb{Q}(\zeta_n)$ be the n-th cyclotomic field, with ring of integers $\mathcal{O}_K$, and let $m = \varphi(n) = \#\{1 \le k < n \mid \gcd(k,n) = 1\}$. Then the following hold:*

1. $[K : \mathbb{Q}] = m$,
2. $\mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q}) \cong (\mathbb{Z}/n\mathbb{Z})^\times \cong (\mathbb{Z}/m\mathbb{Z})$,
3. $\mathcal{O}_K = \mathbb{Z}[\zeta_n]$,
4. *For $n > 2$,* $\Delta_K = (-1)^{m/2}\frac{n^m}{\prod_{p|n} p^{m/(p-1)}}$.

Proof. These are all standard facts. See for example Chapter 2 of the book [Was97]. □

A number field K of degree n over $\mathbb{Q}$ has n embeddings (injective field homomorphisms in $\mathbb{C}$). r_1 are real embeddings, and $2r_2$ are complex embeddings, and they come in conjugate pairs. In total, $n = r_1 + 2r_2$. For example in the cyclotomic case, $\mathbb{Q}(\zeta_n)$ is Galois over K so $r_1 = 0$ and $r_2 = \frac{\varphi(n)}{2}$.

Theorem 3.1.2 (Dirichlet's unit theorem)**.** *The group of units in the ring $\mathcal{O}_K$ of a number field K is finitely generated and has rank $r = r_1 + r_2 - 1$, where r_1 is the number of real embeddings of K and r_2 the number of pairs of conjugate complex embeddings of K.*

Proof. See [Koc00], section 2.9. □

In particular, in the cyclotomic case we get that $\mathbb{Z}[\zeta_n]^\times$ is finitely generated of rank $\frac{\varphi(n)}{2} - 1$.

Definition 3.1.3 (Cyclotomic units)**.** Let $n \not\equiv 2 \bmod 4$ and $K = \mathbb{Q}(\zeta_n)$ be the n-th cyclotomic field. The group of cyclotomic units is defined by

$$C_K = (-1)^{\mathbb{Z}} \zeta_n^{\mathbb{Z}} (1 - \zeta_n)^{\mathbb{Z}} (1 - \zeta_n^2)^{\mathbb{Z}} \dots (1 - \zeta_n^{n-1})^{\mathbb{Z}} \cap \mathcal{O}_K^\times.$$

Cyclotomic units are a chunk of the units of cyclotomic number fields that we understand well. In the prime-power case, they come with satisfying structure.

Lemma 3.1.4. *Suppose $n = p^k$ with p prime and $K = \mathbb{Q}(\zeta_n)$. Then the group C_K of cyclotomic units is generated by (-1), ζ_n, and the units*

$$b_j = \frac{1 - \zeta_n^j}{1 - \zeta_n}, \text{ for } 1 < j < \frac{1}{2} p^k \text{ and } \gcd(j, p) = 1.$$

Proof. See [Was97], Lemma 8.1 for full details. The idea is to explicitly show that real cyclotomic units can be written as a product of powers of (-1) and the b_j's, and then use the fact that in the prime-power case, units can be written as a product of a real units with a n-th root of unity. □

Note that for $k = 1$, in the prime case, C_K is generated by (-1), ζ_p, and the b_j's for $j \in \{2, \dots, \frac{p-1}{2}\}$.

3.2 Ideals

An *ideal* I of $\mathcal{O}_K$ is an additive subgroup of $\mathcal{O}_K$, such that for $r \in \mathcal{O}_K$ and $x \in I$, $rx \in \mathcal{O}_K$. An ideal $\mathfrak{p}$ of $\mathcal{O}_K$ is said to be *prime* if it is not $\mathcal{O}_K$ and if for any $a, b \in \mathcal{O}_K$ such that $ab \in \mathfrak{p}$, $a \in \mathfrak{p}$ or $b \in \mathfrak{p}$. The norm of an ideal $\mathfrak{a}$ is defined as the index $\mathcal{N}(\mathfrak{a}) = \#(O_K/\mathfrak{a})$. $\mathcal{N}(\mathfrak{a})$ is a prime number if and only if $\mathfrak{a}$ is a prime ideal.

A *fractional ideal* of $\mathcal{O}_K$ is a set $J \subset K$ such that there exists a non-zero $r \in \mathcal{O}_K$ such that $rJ \subset$ is an ideal of $\mathcal{O}_K$. In this case we define its norm by $\mathcal{N}(J) = \frac{\mathcal{N}(rJ)}{|\mathcal{N}(r)|}$, and easily check that this generalises the norm for integral ideals.

Proposition 3.2.1. *The set $\mathcal{I}_K$ of fractional ideals of $\mathcal{O}_K$ is an abelian group, with identity $\mathcal{O}_K$.*

A fractional ideal is said to be *principal* if it can be generated by a single element. In this case for an element $g \in K$ we denote by (g) the principal ideal generated by g. The set of principal ideals of $\mathcal{O}_K$ is denoted $\mathcal{P}_K$. In number fields, we have the following theorem that generalises factorisation of rational numbers:

Theorem 3.2.2. *Any fractional ideal $\mathfrak{a} \in \mathcal{I}_K$ can be factored uniquely up to reordering into*

$$\mathfrak{a} = (\mathfrak{p}_1 \dots \mathfrak{p}_r)(\mathfrak{q}_1 \dots \mathfrak{q}_s)^{-1},$$

where $\mathfrak{p}_1, \dots, \mathfrak{p}_r, \mathfrak{q}_1, \mathfrak{q}_s$ are prime (integral) ideals of $\mathcal{O}_K$.

If $u \in \mathcal{O}_K^\times$, the ideal (u) is just $\mathcal{O}_K$.

Definition 3.2.3 (S-units)**.** Let S be a finite set of prime ideals. We say that $s \in K$ is an S-unit if and only if the principal fractional ideal (s) it generates can be written as a product of prime ideals of S using positive or negative powers. Formally, if $S = \{\mathfrak{p}_1, \dots, \mathfrak{p}_k\}$, s is an S-unit if and only if there exists exponents $(e_1, \dots, e_k) \in \mathbb{Z}^k$ such that $(s) = \prod_{i=1}^k \mathfrak{p}_i^{e_i}$.

The set of S-units forms a group that we note $\mathcal{O}_{K,S}^\times$, and by Dirichlet's unit theorem, it is finitely generated of rank $r + k$, where r is the rank of $\mathcal{O}_K^\times$ and k is the number of elements in S.

3.3 Lattices in Number Fields

Ideals as Lattices Recall that a number field of degree n over $\mathbb{Q}$ comes with n embeddings: r_1 real and r_2 pairs of complex embeddings. In the order we write them $\sigma_1, \ldots, \sigma_n$. If $K = \mathbb{Q}[X]/\langle P \rangle$ with P a degree n irreducible polynomial, then any element of K can be written as a linear combination of $(1, X, \ldots, X^{n-1})$. We can always switch from one representation of K to another by going from $\mathbb{Q}$ adjoin a root of the polynomial, to $\mathbb{Q}[X]$ quotiented by the minimal polynomial of the element we adjoin by. The other roots of this polynomial correspond to the images via the embeddings of adjuncted element. This enables us to define two maps from K:

- The *coefficient* embedding:

$$\begin{array}{rccc} \Sigma & : & K & \to & \mathbb{R}^n \\ & & \sum_{i=0}^{n-1} a_i X^i & \mapsto & (a_0, \ldots, a_{n-1}) \end{array};$$

- The *canonical* or *Minkowski* embedding:

$$\begin{array}{rccc} \sigma & : & K & \to & \mathbb{R}^{r_1} \times \mathbb{C}^{r_2} \\ & & x & \mapsto & (\sigma_1(x), \ldots, \sigma_{r_1+r_2}(x)) \end{array};$$

where in the case of σ we can also consider it from K to $\mathbb{R}^{r_1} \times \mathbb{R}^{2r_2}$ by separating real and imaginary parts of the complex embeddings. With this in mind, we have the following proposition, that justifies why we can freely identify ideals to lattices.

Proposition 3.3.1. *Let $\mathfrak{a}$ be a fractional ideal of $\mathcal{O}_K$, then $\Sigma(\mathfrak{a})$ and $\sigma(\mathfrak{a})$ are both full-rank lattices of $\mathbb{R}^n$. Moreover, we have* $\mathrm{vol}(\mathfrak{a}) = \mathcal{N}(\sigma(\mathfrak{a}))\sqrt{|\Delta_K|}$.

Proof. See the classic textbook [ST02], Chapter 8 and Theorem 9.4. □

All notations for lattices extend to ideals via the canonical embedding, for example $\lambda_1(\mathfrak{a})$ now denotes the length of the smallest non-zero vector in the lattice $\sigma(\mathfrak{a})$, and $\|x\|$ for $x \in K$ denotes the Euclidean norm of $\sigma(x)$. It is worth explaining why we choose to look at the lattice from σ and not Σ. In fact, both embeddings have a very similar geometry so the distinction is only minor: on the one hand the coefficient embedding is much better for easy and efficient implementation, and on the other hand the canonical embedding is much more natural from a mathematical standpoint.

The following proposition will be useful:

Proposition 3.3.2. *Let $\mathfrak{a}$ be an ideal of $\mathcal{O}_K$ where K is a cyclotomic number field of conductor n, then*

$$\frac{1}{\mathrm{poly}(n)}\mathcal{N}(\mathfrak{a})^{\frac{1}{n}} \leq \lambda_1(\mathfrak{a}) \leq \mathrm{poly}(n)\mathcal{N}(\mathfrak{a})^{\frac{1}{n}}.$$

Proof. Using Minkowski's first theorem, we get that for any lattice $\mathcal{L}$ of dimension n, $\lambda_1(\mathcal{L}) \leq \sqrt{n}(\mathrm{vol}(\mathcal{L}))^{\frac{1}{n}}$. Combining this with the fact that $\mathrm{vol}(\sigma(\mathfrak{a})) = \mathcal{N}(\mathfrak{a})\sqrt{|\Delta_K|}$ as in Proposition 3.3.1 and with point (4) of Proposition 3.1.1 we get the desired upper bound. For the lower bound we use that if $x \in \mathfrak{a}$, (x) is a sublattice of $\mathfrak{a}$ so $\mathcal{N}(\mathfrak{a})$ must divide $|\mathcal{N}(x)|$ meaning that $|\mathcal{N}(x)| \geq \mathcal{N}(\mathfrak{a})$. Therefore

$$\mathcal{N}(\mathfrak{a})^2 \leq \mathcal{N}(x)^2 = \prod_{i=1}^{n} \sigma_i(x)^2 \leq \left(\frac{\sum_{i=1}^{n} \sigma_i(x)^2}{n}\right)^n = \frac{\|x\|^n}{n^n},$$

where we used the inequality between the arithmetic and geometric means. Taking $x \in \mathfrak{a}$ such that $\|x\| = \lambda_1(\mathfrak{a})$ proves the left inequality. □

Auxiliary Lattices Recall from Theorem 3.1.2 that the unit and S-unit groups are finitely generated, therefore their logarithms should form some sort of lattice. The crucial step in the algorithms we aim to describe searches for units (or S-units) close to a given point in K. In order to turn this step into a closest vector problem instance, we need to formally define these lattices. First we define the set of infinite place $S_\infty = \{\sigma_1, \ldots, \sigma_{r_1+r_2}\}$ to be all embeddings of K up to conjugation. $[K_\sigma : \mathbb{R}]$ is 1 if σ is a real embedding, 2 otherwise.

We define the following embeddings, with $k = \#S$:

- The Log-embedding:

$$\begin{array}{rccc} \mathrm{Log} & : & K & \to & \mathbb{R}^{r_1+r_2} \\ & & \alpha & \mapsto & ([K_\sigma : \mathbb{R}] \cdot \ln|\sigma(\alpha)|)_{\sigma \in S_\infty} \end{array};$$

- The Log-S-embedding:

$$\begin{array}{rccc} \mathrm{Log}_S & : & K & \to & \mathbb{R}^{r_1+r_2+k} \\ & & \alpha & \mapsto & (\mathrm{Log}(\alpha), \{-\mathrm{ord}_{\mathfrak{p}}(\alpha) \cdot \ln(\mathcal{N}(\mathfrak{p}))\}_{\mathfrak{p} \in S}) \end{array};$$

  where $\mathrm{ord}_{\mathfrak{p}}(\alpha)$ is the number of times (positive or negative) that $\mathfrak{p}$ divides (α).

In the first case, we can see that for $\alpha \in K$, the sum Σ of the coordinates of $\mathrm{Log}(\alpha)$ is

$$\begin{aligned} \Sigma &= \sum_{i=1}^{r_1} \ln|\sigma_i(\alpha)| + \sum_{j=1}^{r_2} 2\ln|\sigma_j(\alpha)| \\ &= \ln|\mathcal{N}(\alpha)|, \end{aligned}$$

so that the image $\mathrm{Log}(\mathcal{O}_K^\times)$ of all elements with norm 1 is contained in the hyperplane of $\mathbb{R}^{r_1+r_2}$ defined by

$$H_0^{r_1+r_2} = \{(x_1, \ldots, x_{r_1+r_2}) \in \mathbb{R}^{r_1+r_2} \mid x_1 + \ldots + x_{r_1+r_2} = 0\}.$$

Recall from Dirichlet's unit theorem that $\mathcal{O}_K^\times$ is finitely generated of rank $r_1 + r_2 - 1$, therefore, $\Lambda = \mathrm{Log}(\mathcal{O}_K^\times)$ is in fact a lattice of $\mathbb{R}^{r_1+r_2}$, that spans exactly $H_0^{r_1+r_2}$ when seen as an $\mathbb{R}$-vector space. This lattice Λ is called the *log-unit lattice*.

In a similar fashion, if $\alpha = \prod_{\mathfrak{p}_i \in S} \mathfrak{p}_i^{e_i}$, then the sum Σ_S of the coordinates of $\mathrm{Log}_S(\alpha)$ is

$$\begin{aligned}\Sigma_S &= \sum_{i=1}^{r_1+r_2} [K_{\sigma_i} : \mathbb{R}] \ln |\sigma_i(\alpha)| + \sum_{j=1}^{k} (-\operatorname{ord}_{\mathfrak{p}_j}(\alpha) \ln(\mathcal{N}(\mathfrak{p}_j)) \\ &= \ln |\mathcal{N}(\alpha)| - \ln \left(\prod_{\mathfrak{p} \in S} \mathcal{N}(\mathfrak{p})^{\operatorname{ord}_{\mathfrak{p}}(\alpha)} \right) \\ &= 0,\end{aligned}$$

so $\mathrm{Log}_S(\mathcal{O}_{K,S}^\times)$ is a lattice of $\mathbb{R}^{r_1+r_2+k}$ of rank r_1+r_2+k-1, contained in and spanning $H_0^{r_1+r_2+k}$. This lattice Λ_S is called the *log-S-unit lattice*.

3.4 Class Groups and the Stickelberger lattice

The *class group* $\mathrm{Cl}_K = \mathcal{I}_K/\mathcal{P}_K$ is the quotient of the group of fractional ideals of $\mathcal{O}_K$ by the subgroup of principal fractional ideals of $\mathcal{O}_K$. Two (fractional) ideals $\mathfrak{a}$ and $\mathfrak{b}$ are equivalent if they have the same class. The class of $\mathcal{O}_K$ is clearly the unit in Cl_K, and any ideal in the same class as $\mathcal{O}_K$ is principal. In a way, Cl_K encompasses how principal (or not) an ideal really is. The class group is finite, and its order $h_K = \#\,\mathrm{Cl}_K$ is called the *class number*. In the cyclotomic case, bounds on the size of h_K are known.

Lemma 3.4.1. *Let K be the n-th cyclotomic number field, then*

$$\log h_K = \Theta(n \log \varphi(n)).$$

Proof. See [Was97], Theorem 4.20. □

The maximal real subfield of K is denoted K^+. The class number of Cl_{K^+} is denoted h^+, and will play an important role later. In fact, we have a canonical surjective morphism from Cl_K into Cl_{K^+} obtained by multiplication of ideals by their complex conjugate. The kernel of this morphism defines the *minus* part of the class group, Cl_K^-. Its cardinality is h^-, and clearly we have $h = h^+h^-$.

Recall from Proposition 3.1.1 that $\mathrm{Gal}(\mathbb{Q}(\zeta_n)/\mathbb{Q}) \cong (\mathbb{Z}/n\mathbb{Z})^\times$, where the isomorphism sends integers a coprime with n to morphisms σ_a that send ζ_n to ζ_n^a. Let $\{x\} = x - \lfloor x \rfloor$ denote the fractional part of the real number x.

Definition 3.4.2 (Stickelberger Ideal)**.** For an integer $a \in \mathbb{Z}$, we define the Stickelberger elements in K to be

$$\theta(a) = \sum_{b \in (\mathbb{Z}/n\mathbb{Z})^\times} \left\{ \frac{ab}{n} \right\} \sigma_b^{-1} \in \mathbb{Q}[G].$$

Define the Stickelberger ideal as $\mathscr{S} = \mathbb{Z}[G] \cap \theta\mathbb{Z}[G]$, where $\theta := \theta(1)$.

The Galois ring $\mathbb{Z}[G]$ acts on ideals of $\mathcal{O}_K$ as follows: for $\mathfrak{a}$ an ideal of $\mathcal{O}_K$, and $\alpha = \sum_{\sigma \in G} a_\sigma \sigma \in \mathbb{Z}[G]$, we write

$$\mathfrak{a}^\alpha = \prod_{\sigma \in G} (\mathfrak{a}^\sigma)^{a_\sigma},$$

and this induces a group action on Cl_K. The Stickelberger ideal can also be seen as a sublattice of $\mathbb{Z}^{\varphi(n)}$ as the latter is isomorphic to $\mathbb{Z}[G]$. The following property will be extremely useful for finding class relations. Its proof can be skipped on a first read, but it is interesting to understand where the Stickelberger lattice comes from.

Theorem 3.4.3. *Let $\mathfrak{a}$ be a fractional ideal of $\mathcal{O}_K$, and let $\beta \in \mathscr{S}$. Then the ideal $\mathfrak{a}^\beta$ is principal. In other words, the Stickelberger ideal annihilates the ideal class group of K via the action of the Galois ring $\mathbb{Z}[G]$.*

Proof. Even thought this theorem remains true for arbitrary number fields, we only give a proof in the cyclotomic case for $K = \mathbb{Q}(\zeta_n)$. This proof is the same as the proof in [Was97], Chapter 15 (see Chapter 6 for the more general case).
Let $\beta \in \mathbb{Z}[G]$ such that $\beta\theta \in \mathbb{Z}[G]$. We must show that $\beta\theta$ annihilates any ideal class $\mathfrak{C} \in \mathrm{Cl}_K$. By Chebotarev's density theorem, there exists infinitely many prime ideals of degree 1 in $\mathfrak{C}$. Fix λ such a prime ideal, and let ℓ be the rational prime below λ. Since ℓ splits completely in K, as λ is of degree 1, we must have $\ell \equiv 1 \bmod n$ (see for example [Was97], Theorem 2.13). We let $L = \mathbb{Q}(\zeta_\ell)$ and $M = \mathbb{Q}(\zeta_n, \zeta_\ell)$. From Proposition 3.1.1 point (4) we get that ℓ divides the discriminant Δ_L so is totally ramified in L. Therefore because n and ℓ are coprime, λ is totally ramified in M, and there exists an ideal $\mathcal{L}$ of M such that $\lambda = \mathcal{L}^{\ell-1}$. We now take s a primitive root modulo ℓ, the character $\chi : (\mathbb{Z}/\ell\mathbb{Z})^\times \to \mathbb{C}^\times$ defined by $\chi(s) = \zeta_n$, and consider the Gauss sum

$$g(\chi) = \sum_{b=1}^{\ell-1} \chi(b) \zeta_\ell^b \in M.$$

We are interested in the factorisation of the ideal $(g(\chi))$ in M. First see that

$$\begin{aligned} g(\chi)\overline{g(\chi)} &= \sum_{a,b\in(\mathbb{Z}/\ell\mathbb{Z})^\times} \chi(a)\zeta_\ell^a \chi(b^{-1})\zeta_\ell^{-b} \\ &= \sum_{a,b\in(\mathbb{Z}/\ell\mathbb{Z})^\times} \chi(ab^{-1})\zeta_\ell^{a-b} \\ &= \sum_{b,(ab^{-1})\neq 0} \chi((ab^{-1}))\zeta_\ell^{b(ab^{-1})-b} \\ &= \sum_{b\neq 0}\chi(1) + \sum_{(ab^{-1})\neq 0,1} \chi((ab^{-1}))\sum_{b\neq 0}\zeta_\ell^{b(ab^{-1}-1)} \\ &= (\ell-1) + \sum_{(ab^{-1})\neq 0,1}\chi((ab^{-1}))(-1) = \ell, \end{aligned}$$

therefore only primes above ℓ in M can divide $(g(\chi))$. Now from the definition of ℓ, such primes have to of the form $\sigma(\mathcal{L})$ for a $\sigma \in \mathrm{Gal}(M/\mathbb{Q})$, but since $\lambda = \mathcal{L}^{\ell-1}$ these prime ideals only depend on the restriction of σ to the initial Galois group G. Therefore there exists integers r_a for a coprime to n such that

$$(g(\chi)) = \prod_{\gcd(a,n)=1} \left(\sigma_a^{-1}\mathcal{L}\right)^{r_a},$$

where the σ_a span G. From $(\ell) = \prod_{\sigma\in G}(\sigma\mathcal{L})^{\ell-1}$ and $g(\chi)\overline{g(\chi)} = \ell$ we must have $0 \le r_a \le \ell-1$. Elements of $\mathrm{Gal}(M/\mathbb{Q})$ act on $g(\chi)$ as multiplication by an n-th root of unity. Therefore $g(\chi)^n \in K$, and since n divides $\ell-1$, $g(\chi)^{\ell-1} \in K$. But from the equation above and $\lambda = \mathcal{L}^{\ell-1}$, we deduce

$$(g(\chi))^{\ell-1} = \prod_{\gcd(a,n)=1} \left(\sigma_a^{-1}\lambda\right)^{r_a}; \tag{3.1}$$

which is exactly showing that $\sum_{\gcd(a,n)=1} r_a\sigma_a^{-1} \in \mathbb{Z}[G]$ annihilates the class $\mathfrak{C}$ of λ in Cl_K. We are almost done, but we first need to investigate further the values of the exponents r_a, using some modular trickery. Fix a coprime with n and define $\tau \in \mathrm{Gal}(M/K)$ by $\tau : \zeta_\ell \to \zeta_\ell^s$ where s is the primitive root used for defining χ. Then τ acts trivially mod $\sigma_a^{-1}\mathcal{L}$, and using the simple calculation

$$\frac{\zeta_\ell^s - 1}{\zeta_\ell - 1} = 1 + \zeta_\ell + \ldots + \zeta_\ell^{s-1} \equiv s \bmod \sigma_a^{-1}\mathcal{L};$$

we get

$$\begin{aligned} \frac{g(\chi)}{(\zeta_\ell-1)^{r_a}} &\equiv \left(\frac{g(\chi)}{(\zeta_\ell-1)^{r_a}}\right)^\tau \bmod \sigma_a^{-1}\mathcal{L} \\ &\equiv \frac{g(\chi)}{(\zeta_\ell-1)^{r_a}}\cdot\frac{\chi(s)^{-1}}{s^{r_a}} \bmod \sigma_a^{-1}\mathcal{L}, \end{aligned}$$

but $\sigma_a^{-1}\mathcal{L}$ divides $(\zeta_\ell - 1)$ exactly once, so by definition of r_a, $\frac{g(\chi)}{(\zeta_\ell - 1)^{r_a}}$ has no more factors $\sigma_a^{-1}\mathcal{L}$, therefore we can cancel out and get

$$\zeta_n^{-1} = \chi(s) \equiv s^{-r_a} \bmod \sigma_a^{-1}\mathcal{L},$$

and

$$\zeta_n^{-a} = \sigma_a(\zeta_n) \equiv \sigma_a s^{-r_a} \bmod \mathcal{L}.$$

Both sides of the congruence are in K, so it also holds modulo λ. The order of ζ_m modulo λ is m (as from the equation $\ell = \prod_{j=1}^{\ell-1}(1 - \zeta_\ell^j)$ all roots of unity are distinct modulo λ), so if $\zeta_n^{-1} \equiv s^b \bmod \lambda$, then we must have $b = \frac{(\ell-1)c}{n}$ for some integer c coprime to n. Then it also holds that

$$s^{\frac{(\ell-1)ac}{n}} \equiv s^{r_a} \bmod \ell,$$

and so

$$r_a \equiv \frac{(\ell-1)ac}{n} \bmod (\ell - 1),$$

but since $0 \leq r_a \leq \ell - 1$, we must have $r_a = (\ell - 1)\left\{\frac{ac}{n}\right\}$. Therefore from Equation 3.1, we now get that

$$\sum_{\gcd(a,n)=1} r_a \sigma_a^{-1} = \sum_{\gcd(a,n)=1} (\ell - 1)\left\{\frac{ac}{n}\right\}\sigma_a^{-1} = (\ell - 1)\sigma_c\theta$$

annihilates $\mathfrak{C}$, with $(g(\chi))^{\ell-1} = \lambda^{(\ell-1)\sigma_c\theta}$.

We are now almost done, as we aim to prove $\lambda^{\beta\theta} = 1$. Let $\gamma = g(\chi)^{\sigma_c^{-1}\beta}$, so $\gamma^{\ell-1} \in K$ and with ideals

$$\lambda^{\beta\theta(\ell-1)} = (\gamma^{\ell-1}). \tag{3.2}$$

Since we have taken $\beta\theta \in \mathbb{Z}[G]$, $(\gamma^{\ell-1})$ is the $(\ell - 1)$-th power of an ideal in $\mathbb{Q}(\zeta_\ell)$, so by [Was97], Exercise 9.1, the extension $K(\gamma)/K$ can only be ramified at primes dividing $\ell - 1$. However we already know that $K(\gamma) \subseteq M$ and M/K is totally ramified at ℓ. Therefore the extension $K(\gamma)/K$ is trivial and $\gamma \in K$. Taking the $(\ell - 1)$-th root of Equation 3.2 proves that $\lambda^{\beta\gamma}$ is principal, and therefore we are done. □

3.5 Quantum Algorithms for Class Group Computation

The class group Cl_K is a complicated object, and even though it has been studied extensively throughout the past few centuries, computations in the class group are still difficult. The best classical algorithms run in subexponential time (in $\log|\Delta_K|$),

see [BF14], [BEF+17]. Instances of the hidden subgroup problem in quantum computing are often the key to unlocking quantum polynomial-time algorithms for tasks that have no known classically polynomial algorithms, for instance that is the case with integer factoring and Shor's algorithm. In [EHKS14], Eisenträger, Hallgren, Kitaev and Song generalise the hidden subgroup problem from discrete groups to $\mathbb{R}^n$, enabling computation of unit groups in polynomial-time. Biasse and Song [BS16] build upon [EHKS14] to come up with a polynomial-time quantum algorithm to compute S-unit groups over arbitrary number fields, thus enabling fast quantum class group computations and solving the principal ideal problem in quantum polynomial-time:

Theorem 3.5.1. *There is a quantum algorithm for deciding if an ideal $\mathfrak{a} \subset \mathcal{O}$ of an order $\mathcal{O}$ in a number field K is principal, and for computing $\alpha \in \mathcal{O}$ such that $\mathfrak{a} = \alpha\mathcal{O}$ which runs in polynomial time in the parameters $n = \deg(K)$, $\log(\mathcal{N}(\mathfrak{a}))$ and $\log(|\Delta|)$, where Δ is the discriminant of $\mathcal{O}$.*

Proof. This is [BS16], Theorem 1.3. □

From the algorithm to compute S-units in [BS16], it is not too difficult to derive an algorithm for solving the class group discrete logarithm problem:

Theorem 3.5.2. *Let $\mathfrak{B}$ be a set of prime ideals generating Cl_K and $B \in \mathbb{R}$ such that $\mathcal{N}(\mathfrak{p}) \leq B$ for all $\mathfrak{p} \in \mathfrak{B}$. Then there is a quantum algorithm that given an ideal $\mathfrak{a}$ in $\mathcal{O}_K$ outputs a vector $e \in \mathbb{Z}^{\mathfrak{B}}$ such that $\prod_{\mathfrak{p}\in\mathfrak{B}} \mathfrak{p}^{e_\mathfrak{p}} \sim \mathfrak{a}$, which runs in polynomial time in the parameters $n = \deg(K)$, $\log(\mathcal{N}(\mathfrak{a}))$, $\log(B)$ and $\#\mathfrak{B}$.*

Proof. See [CDW17], Appendix B. □

Chapter 4

Lattice-Based Cryptography and Ideal-SVP

In Chapter 2, we saw examples of hard lattice problems such as SVP, CVP, and their approximation counterparts. In this chapter we explore how these problems are used to build cryptographic protocols. Lattice-based protocols are convenient because of their simplicity and efficiency, but more so because of the strong security proofs that come with it. These security proofs come in the form of *worst-case to average-reductions*, where an average-case problem is proven to be at least as hard as any instance of another problem, which is assumed to be difficult. For example, an attacker would typically be interested in an average-case problem, say a random lattice used in an instantiation of the protocol, and this will be at least as hard as solving any instance of the difficult problem in the security reduction, in particular the worst-case or hardest instance of this problem. But why are lattice-based protocols resistant to attacks by quantum computers? The answer is no one knows, but there is hope as no dangerous attacks are known. Time will tell as more cryptanalysis research is needed.

4.1 Lattice-Based Cryptosystems

The main problems that serve as a basis for lattice-based cryptosystems are the *Short Integer Solution* problem introduced in [Ajt96] in 1996 and the *Learning With Errors* problem introduced in [Reg05] in 2005.

Problem 4.1.1 (Short Integer Solution (SIS))**.** *Let $n, m, q, B \in \mathbb{Z}_{>0}$ with $B \ll q$. $\mathbb{Z}_q$ denotes integers modulo q. Given $A \in \mathbb{Z}_q^{n \times m}$ a uniform random matrix, find a non-zero integer vector $z \in \mathbb{Z}^m$ of norm $\|z\| \leq B$ such that $\boxed{A}\,\boxed{z} \equiv 0 \bmod q$.*

Problem 4.1.2 (Learning With Errors (LWE)). *Let $n, m, q, \in \mathbb{Z}_{>0}$. Let $A \in \mathbb{Z}_q^{m \times n}$ be a uniform random matrix, $s \in \mathbb{Z}_q^n$ a random uniform secret vector, and $e \in \mathbb{Z}^n$ a random integer matrix sampled from a discrete Gaussian distribution. Given A and b, where*

$$\boxed{b} := \boxed{A}\ \boxed{s} + \boxed{e} \ mod\ q;$$

recover the secret vector s.

Sampling from a discrete Gaussian distribution can be done with different parameters αq for some $\alpha < 1$, the *error rate*. Essentially, we want b to be short so that decoding is possible. Problem 4.1.2 is called the *Search* variant of LWE. It can be reformulated as a *Decision* variant, where given a pair (A, b) generated either as in the Search variant for a uniform s, or directly from the uniform distribution, one must decide which is the case with non-negligible advantage.
SIS and LWE enjoy average-case to worst-case reductions to difficult problems (see Section 4.3), however for cryptographic applications they require sharing A publicly where A has quadratic size, and this makes them inherently inefficient. Drawing inspiration from the NTRU cryptosystem [HPS98], new schemes were invented, involving more structured matrices, enabling more compact keys and more efficient algorithms.

4.2 Cryptosystems Based on Structured Lattices

The first time structured lattices were explicitly studied in cryptography goes back to 2002 and the Ring-SIS problem, studied in [Mic07]. Similarly in 2009 and 2010, [SSTX09] and [LPR10] introduce a similar adaptation for LWE, called Ring-LWE, which we will describe in more detail.

Problem 4.2.1 (Ring Learning With Errors (Ring-LWE)). *Let $n, m, q \in \mathbb{Z}_{>0}$. Let R be a ring of degree n over $\mathbb{Z}$, and $R_q = R/qR$. Let $\alpha < 1$ be an error rate and χ be the discrete Gaussian distribution over R with parameter αq. Let $a \in R_q^m$ uniformly at random, $s \in R_q$ a random uniform secret, and let $e \in R^m$ by a vector whose every coordinates are sampled from χ. Given a and b, where*

$$\boxed{b} := \boxed{a}\ \boxed{s} + \boxed{e} \ mod\ q;$$

recover the secret s.

Typically, R is the ring of integers of a number field $K = \mathbb{Q}[X]/\Phi(X)$, where $\Phi(X)$ is a cyclotomic polynomial, for example $X^{2^k} + 1$. As in LWE, this is the search version. Similarly, a decision version exists and [LPR10] proves that its security is equivalent for cyclotomic rings.

In 2011, the paper [BGV11] introduces a variant of Ring-LWE called Module-LWE. Its security is studied in [LS12]. Even though we will not study Module-LWE in this dissertation, it is important to mention it as this is the difficult problem that the newly standardised scheme CRYSTALS-Kyber relies on.

Definition 4.2.2 (Module)**.** Let $k \geq 1$, K a field and R a ring. A subset $M \subseteq K^k$ in an R-module if it is closed under addition and multiplication by elements of R.

If K is a number field and $R = \mathcal{O}_K$, a R-module $M \subseteq K^k$ can be represented as $M = \sum_{i=1^k} I_i \cdot b_i$ where the I_i are ideals of R and the b_i are vectors of K^k. If the I_i are non-zero and the b_i linearly independent, M is said to be a module of rank k. Modules can be seen as lattices via the embeddings σ^k and Σ^k as in section 3.3. Note that rank 1 modules are in fact ideals.

Problem 4.2.3 (Module Learning With Errors (Module-LWE))**.** *Let $n, m, k, q \in \mathbb{Z}_{>0}$. Let R be an ring of degree n over $\mathbb{Z}$. Let $R_q = R/qR$. Let $\alpha < 1$ be an error rate and χ be the discrete Gaussian distribution over R with parameter αq. Let $A \in R_q^{m \times k}$ uniformly at random, $s \in R_q^k$ a random uniform secret vector, and let $e \in R^m$ by a vector whose every coordinates are sampled from χ. Given A and b, where*

$$\boxed{b} := \boxed{A}\ \boxed{s} + \boxed{e} \ mod\ q;$$

recover the secret vector s.

Module-LWE for $k = 1$ is exactly Ring-LWE. For $k > 1$, it is not immediate to see why we use the word module to describe the problem: it is just a general version of Ring-LWE. The module terminology becomes important when studying the underlying lattice and proving security reductions. We will not discuss this further in this work.

4.3 Security Reductions

In this section we give examples of average-case to worse-case reductions for the lattice-based protocols presented earlier in this chapter. The chain of reductions to remember is presented in Figure 4.1.

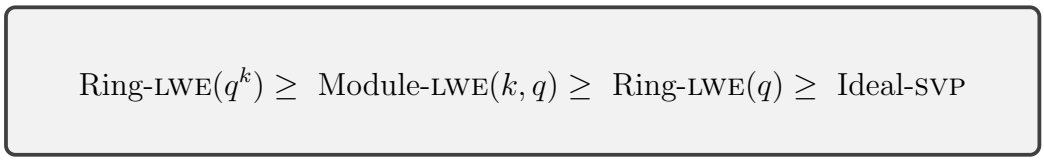

Figure 4.1: Chain of reductions presented in Chapter 4

Theorem 4.3.1. *With the same notations as in Problem 4.1.1 and $m \geq n \log q$, $m = \text{poly}(n)$, solving* SIS *is at least as hard as the worst-case of γ-*SVP *with an approximation factor $\gamma = \text{poly}(n)$.*

Proof. The idea behind this reduction is that $\mathcal{L} = \{z \in \mathbb{Z}^m \mid Az \equiv 0 \bmod q\}$ is a lattice and SIS consists in looking for a short vector in $\mathcal{L}$. See [Ajt96] for full details. □

Theorem 4.3.2. *Using the same notations as in Problem 4.1.2, for $m = \text{poly}(n)$, $q \leq 2^{\text{poly}(n)}$ and an error rate $\alpha \geq 2\frac{\sqrt{n}}{q}$, solving* LWE *is at least as hard as quantumly solving the worst-case of γ-*SVP *with an approximation factor $\gamma = \text{poly}(n/\alpha)$.*

Proof. This reduction relies on classical and quantum computation, and uses the lattice $\mathcal{L} = \{z \in \mathbb{Z}^n \mid \exists s \in \mathbb{Z}^n, As \equiv x \bmod q\}$. For full details see [Reg05]. □

For structured lattices, we have similar reductions. In fact, we are interested in reductions to the following problem: Ideal-γ-SVP (or Ideal-SVP) for short.

Problem 4.3.3 (Ideal Approximate Shortest Vector Problem (Ideal-γ-SVP))**.** *Given an ideal $\mathfrak{a}$ in the ring of integers $\mathcal{O}_K$ of a number field K, and an approximation factor $\gamma \geq 1$, find a non-zero vector $u \in \sigma(\mathfrak{a})$ such that $\|u\| \leq \gamma\lambda_1(\mathfrak{a})$.*

Recall that σ is the Minkowski embedding defined in section 3.3, and that $\lambda_1(\mathfrak{a})$ denotes the length of the shortest non-zero vector in the lattice $\sigma(\mathfrak{a})$. Ideal-γ-SVP is a special case of γ-SVP. It could also have been defined using the coefficient embedding Σ, but this ultimately will not impact the magnitude of γ.

Theorem 4.3.4 ([LPR10])**.** *Using notations of Problem 4.2.1, for any $m = \text{poly}(n)$, ring R of degree n over $\mathbb{Z}$, solving Ring-*LWE *is at least as hard as quantumly solving Ideal-γ-*SVP *on worst-case ideal lattices in R, for some $\gamma = \text{poly}(n)/\alpha$.*

This result means that breaking Ring-LWE automatically solves Ideal-SVP for a polynomial approximation factor. However, this does not necessarily mean that breaking Ideal-SVP breaks Ring-LWE. Their security could be the same, or there could

be a big gap in hardness. The scientific approach is to try and break the weakest link, that is try to find an attack against Ideal-SVP with a polynomial approximation factor. Even if this does not immediately break Ring-LWE, it may provide some insight as to how to break it. The following chapters discuss state-of-the-art attacks against Ideal-SVP. For completeness, we give one last reduction, to put into perspective the security of NIST-standardised scheme CRYSTALS-Kyber.

Theorem 4.3.5 ([AD17], informal)**.** *There is a classical reduction from Module-*LWE *with rank k over a general ring R/qR to Ring-*LWE *in R/q^kR.*

CRYSTALS-Kyber relies on rank $k = 2$ Module-LWE. This last reduction is here to justify that even though there probably is a hardness gap between Module-LWE and Ring-LWE, this gap is still fathomable, as an extremely fast attack on Ring-LWE would enable an attack on Module-LWE. For now though, we stay focused on Ideal-SVP.

Chapter 5

Attacks on Ideal-SVP: An Overview

5.1 Additive Attacks

Until 2014 and [CGS14], most people believed that the best algorithms for solving Ideal-SVP were not better than those for solving SVP on unstructured lattices. These algorithms rely only on the additive structure of the lattice, i.e. they don't use the fact that the lattice was constructed from an ideal with lots of number theoretic structure. Recall from Algorithm 1 that LLL can solve Ideal-SVP with exponential approximation factor in polynomial time. In fact it is also possible to get a polynomial approximation factor in exponential time. The best known trade-offs are listed as Schnorr's hierarchy in [Sch87]. The best variant is the Blockwise Korkine-Zolotarev (BKZ) algorithm [CN11]. It can be adjusted in a way that on lattices of dimension n, it obtain an approximation factor $\exp(\tilde{O}(n^{\alpha}))$ in time $\exp(\tilde{O}(n^{1-\alpha}))$ for any $\alpha \in [0,1]$. The notation $\tilde{O}$ ignores any $\log(n)$ factors, such that $\tilde{O}(n^{\alpha}) = O(n^{\alpha}(\log(n))^{\beta})$ for any $\alpha > 0, \beta \in \mathbb{R}$. Since then, faster algorithms have been discovered that use the multiplicative structure of the ideal, proving that the gap between the hardness of SVP and Ideal-SVP exists.

5.2 Multiplicative Attacks

S-unit attacks, also referred to as unit attacks if S is empty exploit not only the additive structure of algebraic lattices, but also their multiplicative structure. Recall that we are interested in solving Ideal-SVP with approximation factor γ: we are given an ideal $\mathfrak{a} \subset \mathcal{O}_K$ and are asked to output an element $v \in \mathfrak{a}$ of length at most $\gamma\lambda_1(\mathfrak{a})$. In Figure 5.1, we give the general outline of unit attacks.

Unit attack:

1. Start with an ideal $\mathfrak{a}$;
2. Find an ideal $\mathfrak{b}$ of small norm such that $\mathfrak{c} = \mathfrak{a}\mathfrak{b}$ is principal;
3. Find a generator g of $\mathfrak{c}$;
4. In the log-unit lattice, find a vector $\mathrm{Log}(u)$ such that $\mathrm{Log}(g) - \mathrm{Log}(u)$ is short;
5. Output $v = g/u$.

Figure 5.1: High-level description of a unit attack

Phase 3 of the unit attack is solved quantumly by using the algorithm from Theorem 3.5.1. The idea behind Phase 2 is that if we find a short element in $\mathfrak{c}$, the because $\mathfrak{b}$ is small and integral, we get a short element in $\mathfrak{a}$, and we go through the trouble of doing so because Ideal-SVP seems to be easier with principal ideals. Phase 2 is discussed in detail in Section 6.2.

Lemma 5.2.1. *Let g and g' be two generators of a non-zero principal ideal, then $g = g'u$ for a certain $u \in \mathcal{O}_K^\times$.*

Proof. If $g\mathcal{O}_K = g'\mathcal{O}_K$ then there exists $a, b \in \mathcal{O}_K$ such that $g = g'a$ and $g' = gb$. We deduce $g = gba$ and $g' = g'ab$, hence because $g, g' \neq 0$, $ab = ba = 1$ so $a \in \mathcal{O}_K^\times$. □

Suppose that from Phase 3 we know a generator g of $\mathfrak{a}$, and suppose we know that there exists a short generator g' of $\mathfrak{a}$. This is not always the case, but in some early schemes like SOLILOQUY, this is true. Therefore by Lemma 5.2.1, there exists a unit $u \in \mathcal{O}_K^\times$ such that $g = g'u$. By taking the Log, we get $\mathrm{Log}(g') = \mathrm{Log}(g) - \mathrm{Log}(u)$, where $\|\mathrm{Log}(g')\|$ is small and $\mathrm{Log}(u) \in \Lambda$, the log-unit lattice. Therefore Phase 4 can be seen as an instance of CVP in Λ, with target $\mathrm{Log}(g)$. This step can be done efficiently (see Section 6.1) and therefore breaks SOLILOQUY.
In general principal ideals do not have very short generators, this is why it can be useful to use S-unit attacks. S-unit attacks follow the same general pattern, but allow for more choice in $\mathfrak{c}$, leading to maybe better results in the last step of the attack. In Figure 5.2, we give the blueprint for S-unit attacks.

The coefficients α_i are positive, therefore $\mathfrak{c} \subseteq \mathfrak{a}$ as in the unit attack. The hope of such an attack is the the close principal multiple $\mathfrak{c}$ admits an unusually short

S-Unit attack:

1. Start with an ideal $\mathfrak{a}$, and $S = \{\mathfrak{p}_1, \ldots, \mathfrak{p}_k\}$ a finite set of prime ideals;
2. Find positive integer coefficients α_i such that the ideal $\mathfrak{c} = \mathfrak{a}\prod_{i=1}^{k} \mathfrak{p}_i^{\alpha_i}$ is principal;
3. Find a generator g of $\mathfrak{c}$;
4. In the log-S-unit lattice, find a vector $\mathrm{Log}(s)$ such that $\mathrm{Log}(g) - \mathrm{Log}(s)$ is short;
5. Output $v = g/s$.

Figure 5.2: High-level description of an S-unit attack

generator. Adding more primes in S means that there are more candidate $\mathfrak{c}$, so we can hope to find such a good $\mathfrak{c}$. Phase 2 and Phase 3 have polynomial time quantum algorithms from Theorem 3.5.2 and Theorem 3.5.1 respectively. So far there is no guarantee that the α_i are small, or that the ideal $\mathfrak{c}$ has a short generator. This is where the log-S-unit lattice Λ_S comes into play. Recall that the Log_S-embedding sends g to the concatenation of $\mathrm{Log}(g)$ with $(\alpha_i \mathcal{N}(\mathfrak{p}_i))_{i\in\{1,\ldots,k\}}$. If simultaneously we had $\mathrm{Log}(g)$ short and $(\alpha_i \mathcal{N}(\mathfrak{p}_i))_{i\in\{1,\ldots,k\}}$ short, this would mean that the ideal (g) would have a short generator, and that $\prod_{i=1}^{k} \mathfrak{p}_i^{\alpha_i}$ would also be of small norm, meaning the short generator for (g) would end up being a good candidate solution for γ-SVP in $\mathfrak{a}$. Now note that adding a vector from Λ_S to $\mathrm{Log}_S(g)$ has the effect of multiplying (g) with a principal ideal, meaning the resulting ideal is still principal. Therefore Phase 4 as a CVP instance in Λ_S is justified. In Section 6.3 we review how this has been done in the literature. Finally in Phase 5 we must make sure that our output v is still in $\mathfrak{a}$. This is not a real problem and can be made sure of by carefully executing the algorithm for CVP in Phase 4.

5.3 Summary

In Figure 5.3 $\alpha \in [0, 1]$ and $\beta \in [0, \frac{1}{2}]$. All runtime complexities of unit and S-unit attacks are quantum, as they all rely on the algorithms from [BS16] presented in Section 3.5. Using results from [BF14] and [BEF+17], it is possible to do the class

Paper	**Type**	**Number fields**	**Ideals**	γ	**Runtime**	**Pre-p**
[Sch87]	Additive	Any	Any	$2^{\tilde{O}(n^{\alpha})}$	$2^{\tilde{O}(n^{1-\alpha})}$	No
[CDPR16]	Unit	p^k-Cyclotomics	Principal	$2^{\tilde{O}(\sqrt{n})}$	$\text{poly}(n)$	No
[CDW17]	Unit	p^k-Cyclotomics	Any	$2^{\tilde{O}(\sqrt{n})}$	$\text{poly}(n)$	No
[PMHS19]	S-Unit	Any	Any	$2^{\tilde{O}(n^{\beta})}$	$2^{\tilde{O}(n^{1-2\beta})}$	$2^{\tilde{O}(n)}$
[CDW21]	Unit	Cyclotomics	Any	$2^{\tilde{O}(\sqrt{n})}$	$\text{poly}(n)$	No
Conjecture	S-Unit	2^k-Cyclotomics	Any	$\text{poly}(n)$	$2^{\tilde{O}(\sqrt{n})}$	$2^{\tilde{O}(\sqrt{n})}$

Figure 5.3: A summary of recent algorithms for Ideal-γ-SVP

group computations classically, and this would add a term $2^{\tilde{O}(n^{2/3})}$ (or $2^{\tilde{O}(\sqrt{n})}$ if K is prime-power cyclotomic) to all runtimes.

In 2019, Ducas Plançon and Wesolowski in [DPW19] study runtimes behind the unit attack from [CDPR16] and [CDW17] with more precision and run experiments assuming state-of-the-art CVP solvers, and they predict that this multiplicative attack only beats BKZ for cyclotomics of order at least 20000. Of course some improvements are possible, but they give a lower bound on the speed of the attack, arguing with the Gaussian heuristic that the attacks will never be better than BKZ for cyclotomics of order less than 4000. NIST schemes all use fields of order 1024 or less.

The last row of the table refers to conjectured subexponential scalability, announced by Bernstein in a talk at the SIAM conference in 2021 [Ber21], and again by Lange in a talk at the ANTS-XV conference in 2022 [Lan22]. We will take a deeper look at this conjecture in Section 7.3.

Chapter 6

Digging Deeper into Multiplicative Attacks

6.1 Finding Close Vectors in Auxiliary Lattices

In this section, we explore Phase 4 of the unit attack, solving CVP in the auxiliary lattice.

Theorem 6.1.1 ([CDPR16], Theorem 6.5)**.** *Let K be a prime-power cyclotomic number field of conductor n. Then there exists a quantum polynomial-time algorithm that solves Ideal-γ-SVP for any principal ideal $\mathfrak{a}$ of $\mathcal{O}_K$, where γ is an approximation factor of size $\exp(\tilde{O}(\sqrt{n}))$.*

The first step of this algorithm is Phase 2 of the unit attack, for which the algorithm from Theorem 3.5.1 outputs a generator g of the principal ideal $\mathfrak{a}$ in quantum polynomial time. In this section we will provide a rough sketch of the proof of Theorem 6.1.1 in the case where g follows a natural distribution, for example as in SOLILOQUY.

As seen in the overview of unit attacks in Section 5.2, Phase 3 requires that we solve an instance of CVP in the log-unit lattice. Luckily, using cyclotomic units we know an explicit (almost) basis of the log-unit lattice.

Theorem 6.1.2. *Let p be a prime number and $m \in \mathbb{Z}_{>0}$, then group of cyclotomic units C_{p^m} of $\mathbb{Q}(\zeta_{p^m})$ has finite index in the group of units $\mathcal{O}_K^\times$, and more precisely we have*

$$[\mathcal{O}_K^\times : C_{p^m}] = h_{p^m}^+,$$

where $h_{p^m}^+$ is the class number of the maximal real subfield of K.

Proof. The proof follows from combining [Was97], Theorem 8.2 with the result of Exercise 8.5 in the same book. □

From the above theorem, we get that $\mathrm{Log}(C_{p^m})$ is a sublattice of Λ of index $h^+_{p^m}$. Recall Lemma 3.1.4: if $n = p^m$, the group C_{p^m} is generated by (-1), ζ_n, and $b_j = \frac{1-\zeta_n^j}{1-\zeta_n}$ for $1 < j < \frac{n}{2}$ and $\gcd(j,p) = 1$. Let G be the set of such j.

Lemma 6.1.3. *The family of vectors* $\mathrm{Log}(b_j)$ *for* $j \in G$ *is a family of linearly independent vectors in* $H_0^{\frac{\varphi(n)}{2}}$.

Proof. By Lemma 3.1.4, the b_j for $j \in G$ with (-1) and ζ_n generate C_{p^m} the group of cyclotomic units, but by Theorem 6.1.2, $\mathrm{Log}(C_{p^m})$ has the same rank as Λ so $\frac{\varphi(n)}{2} - 1$, therefore as $\mathrm{Log}((-1)^{\mathbb{Z}}\zeta_n^{\mathbb{Z}}) = 0$, we must have that the $\mathrm{Log}(b_j)$ for $j \in G\}$ are a basis of $\mathrm{Log}(C_{p^m})$, in particular they are linearly independent. □

In order to solve CVP in $\mathrm{Log}(C_{p^m})$ we use Babai's round-off algorithm with the basis given by the $\mathrm{Log}(b_j)$. Note that in the general case $h^+_{p^m}$ is not always 1, so we should rather solve CVP in Λ. This is not a big problem and will be discussed in Section 6.4. For now assume all units are cyclotomic units. For the round-off algorithm to work, using Proposition 2.2.12 we would need $\frac{1}{2} \leq \langle \mathrm{Log}(b_j)^\vee, \mathrm{Log}(g) \rangle < \frac{1}{2}$ for all $j \in G$.

The following theorem is again from [CDPR16]. It is quite technical and uses the theory of Dirichlet L-functions to study the geometry of our basis of the log-unit lattice.

Theorem 6.1.4 ([CDPR16], Theorem 3.1)**.** *Vectors of the basis dual to* $(\mathrm{Log}(b_j))_{j\in G}$ *in* $\mathrm{Log}(C_{p^m})$ *all have the same norm and satisfy*

$$\| \mathrm{Log}(b_j)^\vee \|^2 = O(n^{-1} \log^3 n)$$

for all $j \in G$.

Theorem 6.1.5 ([CDPR16], Theorem 4.1)**.** *Assume* c *is an absolute constant and* g *follows a distribution* D *such that for any orthonormal vectors* $v_1, \ldots, v_{\varphi(n)/2-1} \in H_0^{\varphi(n)/2}$*, with probability at least* $\alpha > 0$ $|\langle v_i, \mathrm{Log}(g) \rangle| < c\sqrt{n} \log^{-3/2} n$ *holds for all* i*. Then the unit attack succeeds with probability at least* α.

Proof. We use the round-off algorithm on $\mathrm{Log}(g') = \mathrm{Log}(g) + \mathrm{Log}(u)$, with vectors $(\mathrm{Log}(b_j))_{j\in G}$ as basis. By Theorem 6.1.4 combined with Proposition 2.2.12, the output is $\mathrm{Log}(u)$. Now writing $\mathrm{Log}(u) = \sum_{j\in G} a_j \mathrm{Log}(b_j)$ and outputting $\frac{g}{\prod_{j\in G} b_j^{a_j}}$ gives a shorter vector in the ideal. □

Section 5 of [CDPR16] proves that discrete Gaussian and Gaussian distributions give the assumption in Theorem 6.1.5. But this assumption is in fact not needed if we are ready to accept a larger approximation factor. Section 6 of [CDPR16] proves that a process very similar to this one yields a solution to principal Ideal-SVP with approximation factor $\exp(\tilde{O}(\sqrt{n}))$. It also proves that generators in principal ideals are with high probability of size $\exp(\tilde{O}(\sqrt{n}))$ times that of their shortest vector meaning that algorithms that output a generator can never do better than $\exp(\tilde{O}(\sqrt{n}))$.

6.2 Reducing to Principal Ideals

This section is almost entirely based on the 2017 paper by Cramer, Ducas and Wesolowski [CDW17]. It tackles Phase 2 of the description of unit attacks for prime-power cyclotomics. Theorem 6.1.1 only solves γ-Ideal-SVP in the case of principal ideal. However from Lemma 3.4.1, we expect the class number h_K to be much larger than 1, meaning that Theorem 6.1.1 is not expected to work for general ideals. [CDW17] proposes the following extension:

Theorem 6.2.1 ([CDW17], Main result)**.** *Let K be a prime-power cyclotomic number field of conductor n. Then there exists a quantum polynomial-time algorithm that solves Ideal-γ-SVP for any ideal $\mathfrak{a}$ of $\mathcal{O}_K$, where γ is an approximation factor of size $\exp(\tilde{O}(\sqrt{n}))$.*

This theorem acts as an extension of Theorem 6.1.1, and its proof relies on the following problem (Phase 2 of the unit attack):

Problem 6.2.2 (Close principal multiple problem (CPM))**.** *Given $\mathfrak{a}$ an ideal of $\mathcal{O}_K$, find an ideal $\mathfrak{b}$ of norm less than $\exp(\tilde{O}(n^{3/2}))$ such that $\mathfrak{c} = \mathfrak{a}\mathfrak{b}$ is principal.*

Proposition 6.2.3. *Solving Problem 6.2.2 in quantum polynomial time is enough to prove Theorem 6.2.1.*

Proof. Suppose we have a $\mathfrak{c} = \mathfrak{a}\mathfrak{b}$ principal with $\mathcal{N}(\mathfrak{b}) = \exp(\tilde{O}(n^{3/2}))$, then we can use algorithm 6.1.1 on $\mathfrak{c}$ to get a generator $g \in \mathfrak{c}$ such that

$$\begin{aligned}
\|g\| &\leq \lambda_1(\mathfrak{c}) \exp(\tilde{O}(n^{1/2})) \\
&\leq \mathcal{N}(\mathfrak{c})^{\frac{1}{n}} \operatorname{poly}(n) \exp(\tilde{O}(n^{1/2})) \\
&\leq \mathcal{N}(\mathfrak{a})^{\frac{1}{n}} \mathcal{N}(\mathfrak{b})^{\frac{1}{n}} \operatorname{poly}(n) \exp(\tilde{O}(n^{1/2})) \\
&\leq \mathcal{N}(\mathfrak{a})^{\frac{1}{n}} \exp(\tilde{O}(n^{3/2})) \\
&\leq \lambda_1(\mathfrak{a}) \operatorname{poly}(n) \exp(\tilde{O}(n^{3/2})) \\
&\leq \lambda_1(\mathfrak{a}) \exp(\tilde{O}(n^{3/2})),
\end{aligned}$$

where we used Proposition 3.3.2. $\mathfrak{b}$ is integral, therefore $g \in \mathfrak{c} \subset \mathfrak{a}$ which concludes. □

How hard is it to find a solution to CPM? If the class group is trivial, then so is CPM. If $h_K = \text{poly}(n)$ then we could pick random small norm ideals $\mathfrak{b}$ until $\mathfrak{a} \sim \mathfrak{b}^{-1}$. But h_K is larger, so this won't be enough.

As a first step towards CPM, suppose we have a finite set $\mathfrak{B}$ of polynomially bounded prime ideals that generate Cl_K. We can use Theorem 3.5.2 on the ideal $\mathfrak{a}$ to find in quantum polynomial time a vector $e \in \mathbb{Z}^{\mathfrak{B}}$ such that the relation

$$\mathfrak{a} \sim \prod_{\mathfrak{p} \in \mathfrak{B}} \mathfrak{p}^{e_\mathfrak{p}}$$

holds in Cl_K. If $\mathfrak{b} = \prod_{\mathfrak{p} \in \mathfrak{B}} \mathfrak{p}^{-e_\mathfrak{p}}$, then $\mathfrak{ab}$ is principal. Our objective would be to have $\|e\|_1 = \tilde{O}(n^{3/2})$ so that $\mathcal{N}(b) \leq \text{poly}(n)^{\tilde{O}(n^{3/2})} = \exp(\tilde{O}(n^{3/2}))$, using the polynomial bound on the prime ideals of $\mathfrak{B}$. We will also need to make sure that $\mathfrak{b}$ is integral. Assume that our factor basis for $\mathfrak{B}$ is of the form

$$\mathfrak{B} = \{\mathfrak{p}^\sigma \mid \sigma \in G\},$$

where $G = \text{Gal}(K/\mathbb{Q})$ is the Galois group of K. Then by writing $\alpha = \sum_{\sigma \in G} \sigma e_\sigma \in \mathbb{Z}[G]$, we get that $[\mathfrak{a}] = [\mathfrak{p}]^\alpha$, where $[\cdot] \in \text{Cl}_K$ denotes the class of $\cdot$ in Cl_K. Recall from Theorem 3.4.3 that the Stickelberger lattice $\mathscr{S}$ of $\mathbb{Z}[G]$ is such that any $s \in S$ satisfies $\mathfrak{p}^s$ is principal. Therefore, finding if we somehow knew a decent basis for $\mathscr{S}$, we would be able to solve a CVP instance in the lattice of class relations and find a $\beta \in \mathscr{S}$ such that $\|\alpha - \beta\|_1$ is small. The explanation is getting a bit dense so we make a few remarks for clarity.

- Cl_K is supposed to be generated by the finite set $\mathfrak{B}$, so the set

$$\left\{ e \in \mathbb{Z}^{\mathfrak{B}} \mid \prod_{\mathfrak{p} \in \mathfrak{B}} \mathfrak{p}^{e_\mathfrak{p}} \text{ is principal} \right\}$$

  is indeed a lattice.

- We reduced CPM to an instance of CVP in the above lattice, but for a different norm $\|\cdot\|_1$ instead of the usual Euclidean norm. This won't be a problem as we will see later.

- The Stickelberger lattice is not full-rank, so even with a short basis of $\mathscr{S}$ we cannot immediately reduce. This technicality can be solved by working in Cl_K^-.

Lemma 6.2.4. *Assuming the Generalised Riemann Hypothesis (GRH) and* $h_K^+ = \tilde{O}(\text{poly}(n))$*, there is a polynomial algorithm in* n *and* $\log(\mathcal{N}(\mathfrak{a}))$ *that send a principal ideal* $\mathfrak{a}$ *in* $\mathcal{O}_K$ *to a principal ideal* $\mathfrak{a}' = \mathfrak{a}\mathfrak{b}$ *in* Cl_K^- *such that* $\mathcal{N}(\mathfrak{b}) \leq \exp(\tilde{O}(n))$.

Proof. This will not be discussed here, see [CDW17], Section 4.1. □

The rank of the Stickelberger lattice $\mathscr{S}$ is not enough ($\varphi(n)/2 + 1$ and it should be $\varphi(n)$), so we *augment* so that it has full rank in $\mathbb{Z}[G]$, and still annihilates Cl_K^-.

Definition 6.2.5. The augmented Stickelberger ideal (or lattice) is defined as

$$\mathscr{S}' = \mathscr{S} + (1+\tau)\mathbb{Z}[G],$$

where τ acts on Cl_K as complex conjugation.

The enables us to derive the following useful proposition.

Proposition 6.2.6. *The following points concerning the augmented Stickelberger lattice* $\mathscr{S}'$ *are true:*

1. $\mathscr{S}'$ *annihilates* Cl_K^-*;*
2. $\mathscr{S}'$ *is full-rank in* $\mathbb{Z}[G]$*;*
3. *If* $\varphi(n) \geq 3$ *there exists a set* W *such that* $\forall w \in W, \|w\| \leq 2\sqrt{n}$*, and* W *that generates* $\mathscr{S}'$*.*

Proof. First we prove point (1). Clearly $\mathscr{S}$ annihilates Cl_K^- by Theorem 3.4.3, as Cl_K^- is a subgroup of Cl_K. Also, we know that Cl_K^- is the kernel of the map that sends $\mathfrak{a}$ to $\mathfrak{a}\overline{\mathfrak{a}} = \mathfrak{a}^{(1+\tau)}$, meaning $(1+\tau)\mathbb{Z}[G]$ also annihilates Cl_K^-.

For point (2) we need to use Theorem 6.19 from [Was97] that states that the lattice $\mathscr{S}^- = \mathscr{S} \cap (1-\tau)\mathbb{Z}[G]$ has full-rank in $(1-\tau)\mathbb{Z}[G]$, therefore we deduce that the lattice $\mathscr{S}^- + (1-\tau)\mathbb{Z}[G]$ has full-rank in $2\mathbb{Z}[G] \subset (1+\tau)\mathbb{Z}[G] + (1-\tau)\mathbb{Z}[G]$, and then so does $\mathscr{S}'$.

For proof of point (3) we refer to [CDW17], Lemma 4, for which the set W can be written down explicitly. □

Now we can use W to reduce $\alpha \in \mathbb{Z}[G]$, by using Babai's Nearest plane algorithm (recall Algorithm 2).

Lemma 6.2.7. *Suppose* $\varphi(n) \geq 3$*, let* $\alpha \in \mathbb{Z}[G]$*, then there exists an algorithm that finds an element* $\beta \in \mathbb{Z}[G]$ *such that* $\|\beta\|_1 \leq n^{3/2}$*, and* $\alpha - \beta$ *annihilates* Cl_K^-*, which runs in classical polynomial time in* n *and* $\log \|\alpha\|$*.*

Proof. From Proposition 6.2.6, W provides a basis B for $\mathscr{S}'$ and $\mathscr{S}'$ is full-rank in $\mathbb{Z}[G]$. If $B^\star$ denotes the Gram-Schmidt orthogonalisation of B then by Lemma 2.2.13 Babai's nearest plane algorithm applied to α leads to a vector β such that

$$\|\alpha - \beta\|^2 \leq \frac{1}{4}\sum_{i=1}^{\varphi(n)} \|b_i^\star\|^2,$$

and in particular, $\|\alpha - \beta\| \leq \frac{1}{2}\sqrt{n}\max(\|b_i^\star\|)$, but $\max(\|b_i^\star\|) \leq \max(\|w\|) \leq 2\sqrt{n}$ by property of the orthogonalisation and by point (3) of Proposition 6.2.6. Therefore by Cauchy-Schwarz and combining the previous inequalities,

$$\|\alpha - \beta\|_1 \leq n^{3/2}.$$

The second condition falls from point (2) of the same proposition. □

Now we can give an algorithm that proves Theorem 6.2.1. Start with an ideal $\mathfrak{a}$ of $\mathcal{O}_K$. Suppose Cl_K^- is generated by $\mathfrak{B} = \{\mathfrak{p}^\sigma \mid \sigma \in G\}$, where $\mathcal{N}(\mathfrak{p}) = \tilde{O}(\mathrm{poly}(n))$. Assume GRH and use Lemma 6.2.4 so that we get a principal $\mathfrak{a}'$ in Cl_K^-, with norm close to the norm of $\mathfrak{a}$. Then use the quantum algorithm from Theorem 3.5.2 to get a vector $e \in \mathbb{Z}^{\mathfrak{B}}$ that is a solution of the class group discrete logarithm problem for $\mathfrak{a}'$. Define $\alpha = \sum_{\sigma\in G} \sigma e_\sigma \in \mathbb{Z}[G]$, and use the algorithm from Lemma 6.2.7 to get a $\beta \in \mathbb{Z}[G]$ such that $\|\beta\|_1 = \tilde{O}(n^{3/2})$ and acting as α on Cl_K^-. If $\beta = \beta^+ - \beta^-$ where β^+ and β^- have positive coefficients in $\mathbb{Z}[G]$, then taking $\mathfrak{b} = \mathfrak{p}^{\beta^- + \tau\beta^+}$ ensures that $\mathfrak{b}$ is integral, $\mathfrak{a}\mathfrak{b}$ is principal, and that $\mathcal{N}(\mathfrak{b}) = \exp(\tilde{O}(n^{3/2}))$. All steps are polynomial or quantum polynomial, so we are done.

Remark 6.2.8. We have assumed GRH, $h_K^+ = \mathrm{poly}(n)$, and that there exists a factor basis $\mathfrak{B}$ of Cl_K^- that looked like the Galois orbit of a single small prime. For a discussion of these assumptions see Section 6.4.

6.3 Ideal-svp with Pre-processing

The previous two sections explain unit attacks. This section is inspired by the paper [PMHS19] by Pellet–Mary, Hanrot and Stehlé in 2019. [PMHS19] marks the first time S-units are used for an attack on Ideal-SVP in the literature. Further improvements include [BRL20], and more recently [BLNRL21], enabling experiments on general cyclotomic fields of degree up to 210.

The attack of [PMHS19] relies on the following algorithm of Laarhoven in [Laa17] for CVP with pre-processing.

Theorem 6.3.1 ([Laa17], Corollaries 2 and 3)**.** *Let $\alpha \in [0, \frac{1}{2}]$. Assuming the Gaussian heuristic for a full-rank lattice $\mathcal{L} \subset \mathbb{R}^n$, there exists an algorithm that takes as pre-processing input $\mathcal{L}$ and a target vector $v \in \mathbb{R}^n$ and as query input a vector $u \in \mathcal{L}$ such that $\|u - v\| \leq O(n^\alpha) \operatorname{dist}(v, \mathcal{L})$, with pre-processing time $2^{\tilde{O}(n)}$ and query time $2^{(\tilde{O}(n^{1-2\alpha}))}$.*

The heuristic assumption will be discussed further in Section 6.4. This algorithm enables the following result, true for general number fields, but presented here for cyclotomic number fields for simplicity of exposition.

Theorem 6.3.2 ([PMHS19], Theorem 5.1)**.** *Let K be the n-th cyclotomic number field, and $\alpha \in [0, \frac{1}{2}]$. Under GRH and the Gaussian heuristic, there exists an algorithm that takes K as pre-processing input, and an ideal $\mathfrak{a}$ of $\mathcal{O}_K$ as query input that solves Ideal-γ-SVP where γ is an approximation factor of size $\exp(\tilde{O}(n^{1-2\alpha}))$, and that runs in time $\exp(\tilde{O}(n^\alpha))$ with pre-processing in time $\exp(\tilde{O}(n))$.*

The key take from Theorem 6.3.2 is that there exists a non-uniform algorithm that outperforms all known additive attacks. This is not the case for unstructured lattices and shows a potential gap between Ideal-γ-SVP and γ-SVP for $\gamma = \text{poly}(n)$ ($\alpha = 0$ in Theorem 6.3.2). We give a short explanation of the proof of Theorem 6.3.2, and refer to [PMHS19] and [BRL20] for the full technical details, especially relatively to parameter choices and twisting the log-S-unit lattice.

The algorithm is inspired by the unit attack of [CDPR16] and [CDW17]. It acknowledges that an algorithm that outputs a generator of an arbitrary principal ideal will have no chance of breaking $\exp(\tilde{O}(\sqrt{n}))$ approximation factors. For this reasons we fix a finite set of prime ideals $S = \{\mathfrak{p}_1, \ldots, \mathfrak{p}_k\}$ that generate Cl_K, and as in Section 6.2, we want to use S to obtain a principal ideal close to $\mathfrak{a}$, effectively solving a CVP instance in the lattice of class relations over S. We then use the quantum algorithm to get a generator of the obtained principal ideal, and then need to solve another CVP instance in the log-unit lattice as in Section 6.1. The idea is to fuse both CVP instances into one, inside a concatenated lattice, the log-S-unit lattice. In the Unit-attack, we knew an almost basis of the log-unit lattice with cyclotomic units, and could use it to solve CVP in the log-unit lattice. However in this case we do not have a satisfying basis[1]. The key here is that we do not need to choose S depending on $\mathfrak{a}$, but only on the number field K. Therefore we can entirely precompute the lattice Λ_S as soon as we know K. This justifies why we can use Theorem 6.3.1 (But this

[1]see [BLNRL21] for more work in this direction in the cyclotomic case, where Stickelberger generators obtained with Jacobi sums are used to derive an explicit basis of the log-S-unit

does not justify that Λ_S behaves according to the Gaussian heuristic, this discussion is left for Chapter 7). Concretely the algorithm goes as follows:

Pre-processing phase: For a given K, we first compute an adequate factor base of prime ideals S. It must contain ideals of bounded norm, generate the class group, and have adequate size. The fact that this is always possible relies on a generalised prime number theorem and a theorem of Bach ([Bac90]). Once S is set, we compute a generating set of S-units, along with their Log_S-embeddings. This generates our lattice Λ_S, which we can pre-process using Theorem 6.3.1's pre-processing algorithm. This takes exponential time, even though the S-unit computation part can be sped up quantumly using [BS16].

Query phase: Suppose that $\delta > 0$ is such that the γ-CVP oracle from Theorem 6.3.1 used on any target v outputs $s \in \Lambda_S$ such that $\|s - v\|_\infty \leq \delta$. We first use Theorems 3.5.2 and 3.5.1 to quantumly solve the class group discrete log problem for $\mathfrak{a}$ with factor base S. This gives a $g \in K$ and integers $\alpha_i \in \mathbb{Z}$ such that

$$(g) = \prod_{\mathfrak{p}_i \in S} \mathfrak{p}_i^{\alpha_i}.$$

Define the target

$$v = (\mathrm{Log}(g), \{\alpha_i + \delta\}_{1 \leq i \leq k}) \in \Lambda_S.$$

We then use the CVP solver to get a $\mathrm{Log}_S(s) \in \Lambda_S$ such that $\|\mathrm{Log}_S(s) - v\|_\infty \leq \delta$. In particular if for $i \in \{1, \ldots, k\}$, $\beta_i = \mathrm{ord}_{\mathfrak{p}_i}(s)$ we get $|\beta_i - \alpha_i + \delta| \leq \delta$, so $0 \leq \alpha_i - \beta_i \leq 2\delta$. But

$$\left(\frac{g}{s}\right) = \mathfrak{a} \prod_{i=1}^{k} \mathfrak{p}_i^{\alpha_i - \beta_i};$$

which means that g/s is an element of $\mathfrak{a}$ because $\alpha_i - \beta_i \geq 0$ and the $\mathfrak{p}_i$ are integral, but also that g/s has small norm:

$$|\mathcal{N}(g/s)|^{\frac{1}{n}} \leq \mathcal{N}(\mathfrak{a})^{\frac{1}{n}} B^{\frac{1}{n}\sum_{i=1}^{k}(\alpha_i - \beta_i)} \leq \mathcal{N}(\mathfrak{a})^{\frac{1}{n}} \exp\left(O\left(\frac{\delta k \log B}{n}\right)\right),$$

where B is a bound on the norm of the elements of S. By [Bac90] we can assume $B = O(\log^2 |\Delta_K|)$. Comparing with Lemma 3.3.1, we can see that careful tweaking of k and managing δ allows for a satisfying approximation. Some shortcuts have been made regarding parameter choice, again see [PMHS19] and [BRL20] for the full picture.

6.4 A Comment on Assumptions

Along the way we have seen that the unit and S-unit attacks rely on assumptions and heuristics. Some have extensive evidence that go their way, and do not have a big impact on neither runtime nor correctness. However some would benefit from further study. In this section we present a list of the heuristics that have been used, where they are needed and if there is sufficient evidence to support them.

Generalised Riemann Hypothesis: GRH is a number-theoretical conjecture on the zeros of Dirichlet L-functions, with many consequences in analytic number theory and beyond. It is very widely considered to be true and enjoys very strong computational evidence. For our purposes, GRH implies the result of Bach [Bac90] that says that Cl_K can be generated by a factor basis of prime ideals of norm less than $O(\log^2 |\Delta_K|)$. This and more recent results is used in both [CDW17] (Lemma 6.2.4) and [PMHS19] (Theorem 6.3.2).

Size of h_K^+: In the proof of Theorem 6.1.1 we use a basis of the cyclotomic units C_K instead of one of the full unit group to decode the log-unit lattice. From Theorem 6.1.2 we know their index in the unit group is equal to h_K^+, the class number of the maximal real subfield of K. To tackle the general case, we can use the polynomial quantum algorithm from [EHKS14] to compute the h_K^+ coset representatives of C_K in $\mathcal{O}_K^\times$, meaning we would only multiply the runtime by a factor h_K^+. In Lemma 6.2.4, we use the assumption that $h_K^+ = \tilde{O}(\mathrm{poly}(n))$ to make sure that we can reduce the CPM problem from an principal ideal in Cl_K to a principal ideal in Cl_K^- in polynomial time. This assumption seems very reasonable for prime-power cyclotomic fields and enjoys very strong theoretical and computational evidence. In fact [BPR] conjecture that for all but finitely many pairs (p, k) where p is a prime and m an integer, we have $h^+_{\mathbb{Q}(\zeta_{p^{m+1}})} = h^+_{\mathbb{Q}(\zeta_{p^m})}$. A specialisation of this conjecture is the Weber class number problem, conjecturing that $h_K^+ = 1$ for power-of-two cyclotomics. These are more often than not used in cryptography, so the assumption that $h_K^+ = \tilde{O}(\mathrm{poly}(n))$ seems very reasonable.

Small Galois factor basis of Cl_K^-: In our discussion of Theorem 6.2.1, we suppose that Cl_K^- can be generated by a factor basis of the form $\mathfrak{B} = \{\mathfrak{p}^\sigma | \sigma \in G\}$ where $\mathfrak{p}$ is a prime ideal of polynomially bounded norm. This can be relaxed to d different Galois orbits of polynomially bounded prime ideals. If $d = \tilde{O}(1)$, this does not impact the resulting Ideal-SVP approximation factor. In fact, randomly choosing primes in Cl_K^-

of norm less than a certain bound $B = \text{poly}(n)$ suffices most of the time. This is made formal in [CDW17], Proposition 1. Testing membership of a prime ideal in Cl_K^- can be done quantumly using Theorem 3.5.1.

Gaussian heuristic: The Gaussian heuristic is used for the Laarhoven CVP with pre-processing algorithm in the 6.3.2 attack. Theorem 6.3.1 on a lattice $\mathcal{L}$ of dimension d requires that there exists a constant $c > 0$ such that the ball $c\lambda_1(\mathcal{L})B_d$ contains at least 2^n points of $\mathcal{L}$, and these points behave as uniformly and independently on the unit sphere. This assumption is a bit much to ask from the log-S-unit lattice. See Chapter 7 and [BL21] for reasons why the log-S-unit lattice does not behave according to the Gaussian heuristic. The Gaussian heuristic is used by [Laa17] to predict reduction of the lattice. It is also used by [DPW19] to give heuristic lower bounds on the speed of the attacks [CDPR16],[CDW17]. The risk here is that the attacks work better than expected by the Gaussian heuristic, because of the apparent non-randomness of the log-unit and log-S-unit lattices.

Other heuristics: A couple other heuristics are used in [BF14] for the classical alternatives to the quantum algorithms of [BS16] and [EHKS14], but we won't discuss them here. [PMHS19] also uses more Gaussian heuristic-like assumptions regarding the distribution of the input and output vectors of the Laarhoven CVP algorithm. See Heuristics 5 and 6 in [PMHS19] for additional details.

Chapter 7

Does the log-S-unit Lattice Behave Like a Random Lattice?

7.1 A Simplified log-S-unit Lattice

How does the log-S-unit lattice behave compared to predictions from the Gaussian Heuristic? In this section we go through the reasoning from [BL21] to demonstrate that spherical models give inaccurate predictions for λ_1 by considering the S-unit lattice Λ_S in the extreme case of rationals. This corresponds to the cyclotomic case of degree 1, where $K = \mathbb{Q}$ and $\mathcal{O}_K = \mathbb{Z}$. Let $y \geq 2$ be our smoothness bound. We look at sets of S-units with S of the form $S = \{p\mathbb{Z} \mid \text{prime } p \leq y\}$. In this setting, S-units are rationals of the form $\frac{a}{b}$ where a, b are y-smooth rational integers.

A basis for Λ_S is given by $d = \pi(y)$ rows of the form $(\log p, 0, \ldots, 0, -\log p, 0, \ldots)$ for $p \leq y$ prime, where $\pi(y)$ counts the number of primes less than or equal to y. If p_d is the largest such prime, L has basis:

$$L = \begin{pmatrix} \log 2 & -\log 2 & 0 & 0 & \ldots & 0 \\ \log 3 & 0 & -\log 3 & 0 & \ldots & 0 \\ \log 5 & 0 & 0 & -\log 5 & \ldots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & 0 \\ \log p_d & 0 & 0 & 0 & \ldots & -\log p_d \end{pmatrix}$$

To compute the volume of Λ_S, we factor out $\prod_{p \leq y} \log p$, and multiply the remaining

$d \times (d+1)$ matrix by its transpose to get the following $d \times d$ matrix:

$$\begin{pmatrix} 1 & -1 & 0 & 0 & \dots & 0 \\ 1 & 0 & -1 & 0 & \dots & 0 \\ 1 & 0 & 0 & -1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & 0 \\ 1 & 0 & 0 & 0 & \dots & -1 \end{pmatrix} \begin{pmatrix} 1 & 1 & 1 & \dots & 1 \\ -1 & 0 & 0 & \dots & 0 \\ 0 & -1 & 0 & \dots & 0 \\ 0 & 0 & -1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & 0 \\ 0 & 0 & 0 & \dots & -1 \end{pmatrix} = \begin{pmatrix} 2 & 1 & 1 & \dots & 1 \\ 1 & 2 & 1 & \dots & 1 \\ 1 & 1 & 2 & \dots & 1 \\ \vdots & \vdots & \vdots & \ddots & 1 \\ 1 & 1 & 1 & \dots & 2 \end{pmatrix}$$

This matrix has eigenvalue 1 with multiplicity $d-1$, and therefore also $d+1$ with multiplicity 1 by looking at the trace. therefore its determinant is the product of the eigenvalues: $d+1$. We deduce that

$$\operatorname{vol}(\Lambda_S) = \det L = (1+\#S)^{1/2} \prod_{p \le y} \log p.$$

We now prove that $(\operatorname{vol}\Lambda_S)^{1/d} = (1+o(1))\log y$ as $y \to \infty$. For this we define f the characteristic function of prime numbers, note that $\pi(x) = \sum_{n \le x} f(n)$ and compute using Abel's transformation:

$$\begin{aligned} \log\left(\prod_{p \le y} \log p\right) &= \sum_{n \le y} f(n) \log\log n \\ &= \log\log 2 + \sum_{3 \le n \le y} f(n) \log\log n \\ &= \log\log 2 + \pi(y)\log\log y - \log\log 2 - \int_2^y \frac{\pi(t)}{t\log t} dt \\ &= \pi(y)\log\log y + O\left(\frac{y}{\log y}\right). \end{aligned}$$

Now $\log\left(\prod_{p \le y} \log p\right) \sim \pi(y)\log\log y$ and both terms go to ∞ as $y \to \infty$ therefore we have $\prod_{p \le y} \log p \sim \exp(\pi(y)\log\log y) = (\log y)^d$ so

$$(\operatorname{vol}\Lambda_S)^{1/d} \sim (1+\#S)^{1/2d} \log y \sim \log y,$$

as $\#S = d+1$. The lattice has rank d as this is the rank of the matrix L. Therefore by Lemma 2.4.6, we expect the length λ_S of the shortest non-zero vector in Λ_S to satisfy as $d \to \infty$

$$\lambda_S \sim \sqrt{\frac{d}{2\pi e}} (\operatorname{vol}\Lambda_S)^{1/d} \sim \sqrt{\frac{y}{2\pi e \log y}} \log y = \sqrt{\frac{y\log y}{2\pi e}},$$

where we used the prime number theorem on d. However all vectors $\mathrm{Log}_S(p)$ for $p \leq y$ are in Λ_S and have length much smaller. In particular $\|\mathrm{Log}_S(2)\| = \sqrt{2}\log 2$, which is constant and far from the order of magnitude of λ_S. This proof is not for general cyclotomics, but it gives an idea as to why the log-S-unit lattice cannot be treated like a random lattice. It contains many short vectors and therefore should enable much faster reductions than in the case of random lattices.

7.2 Non-Randomness of the log-unit Lattice

In the last subsection, we saw that spherical models of S-unit lattices in the rational case were getting worse at predicting short vectors as the smoothness bound y increased. Following the reasoning from [BL21], we study the opposite case where $S = \emptyset$ is empty, so that the S-unit lattice is in fact the unit lattice. [BL21] treats the case of power-of-two cyclotomics. We generalise this approach to all prime-power cyclotomics, i.e. fields $K = \mathbb{Q}(\zeta_n)$, where $n = p^k$ with p a prime number and $k \geq 1$ an integer.

Our objective is to compare the actual size of short vectors in $\Lambda = \mathrm{Log}(\mathcal{O}_K^\times)$ with the size predicted by the Gaussian heuristic, as in Lemma 2.4.6. In order to use the formula in Lemma 2.4.6, we need to compute the volume of the log-unit lattice. A proof of a precise asymptotic in the prime-cyclotomic case is given in [DPW19], Appendix B.

Theorem 7.2.1. *Let $K = \mathbb{Q}(\zeta_n)$ be a cyclotomic field with $n = p^k$ a prime power. Then the volume of the log-unit lattice Λ satisfies*

$$\left(\frac{\mathrm{vol}(\Lambda)}{h_K^+}\right)^{\frac{1}{\varphi(n)/2-1}} \sim \sqrt{n}$$

as $n \to \infty$.

Proof. Details can be found in Appendix B of [DPW19]. They involve a result from [Was97] on an expression of the product Rh_K^+, where R is the *regulator* of the number field K. □

Remark 7.2.2. This is not exactly the formula from [DPW19], as we had to account for a slight difference in the definition of Λ.

We can assume that $h_n^{+\frac{1}{\varphi(n)/2-1}} \sim 1$ (this has been extensively discussed in Section 6.4). Along with Lemma 2.4.6, we get that the predicted length λ of the shortest

non-zero vector in Λ, assuming the Gaussian heuristic, satisfies for prime-powers $n \to \infty$

$$\lambda \sim \sqrt{\frac{\varphi(n)/2-1}{2\pi e}}\,\mathrm{vol}(\Lambda)^{\frac{1}{\varphi(n)/2-1}} \sim \sqrt{\frac{n\varphi(n)}{4\pi e}}. \tag{7.1}$$

For prime-power cyclotomics we have $\frac{n}{2} \leq \varphi(n) \leq n$, therefore we expect $\lambda = \Theta(n)$.

We are interested in comparing the shortest vector prediction λ from the Gaussian heuristic with the length of an actual small vector in Λ. For this we consider a particular unit and estimate its length.

Lemma 7.2.3. *Let $p \neq 3$ be a prime number, $k \in \mathbb{Z}_{>0}$ and $K = \mathbb{Q}(\zeta_{p^k})$ the associated prime-power cyclotomic number field, then if $n = p^k$,*

$$u = 1 + \zeta_n + \zeta_n^{-1}$$

is a unit in $\mathcal{O}_K$.

Proof. We have

$$u = \zeta_n^{-1}\frac{1-\zeta_n^3}{1-\zeta_n} = 1 + \zeta_n + \zeta_n^{p^k-1} \in \mathbb{Z}[\zeta_n],$$

and for a certain integer m such that $3m \equiv 1 \bmod p^k$, that exists because $\gcd(3,p) = 1$:

$$u' = \zeta_n\frac{1-\zeta_n}{1-\zeta_n^3} = \zeta_n\frac{1-\zeta_n^{3m}}{1-\zeta_n^3} = \zeta_n\sum_{j=0}^{m-1}\zeta_n^{3j} \in \mathbb{Z}[\zeta_n].$$

Both are elements of $\mathcal{O}_K = \mathbb{Z}[\zeta_n]$ and $uu' = 1$, so we have proved that $u \in \mathbb{Z}[\zeta_n]^\times$. □

We claim that the vector $\mathrm{Log}(u)$ is a non-zero vector shorter than λ, the prediction from the Gaussian heuristic.

Lemma 7.2.4. *Let $n = p^k$ with $p \neq 3$, and $u = 1 + \zeta_n + \zeta_n^{-1}$ in the cyclotomic field $K = \mathbb{Q}(\zeta_n)$. Then*

$$\frac{1}{4}\sqrt{n} \leq \|\,\mathrm{Log}(u)\| \leq \sqrt{2}\log 3\sqrt{n}.$$

Proof. In $K = \mathbb{Q}(\zeta_n)$, the log-unit lattice is obtained from the $\frac{\varphi(n)}{2}$ embeddings obtained from sending ζ_n to ζ_n^j, for elements of $E = \{1 \leq j \leq \frac{1}{2}p^k | \gcd(j,p) = 1\}$. We have

$$\mathrm{Log}(u) = (2\log|1+\zeta_n+\zeta_n^{-1}|, \ldots, 2\log|1+\zeta_n^{\frac{p^k-1}{2}} + \zeta_n^{-\frac{p^k-1}{2}}|).$$

Calculating the norm,

$$\|\,\mathrm{Log}(u)\|^2 = \sum_{j\in E}(2\log|1+\zeta_n^j+\zeta_n^{-j}|)^2 = \sum_{j\in E}\left(2\log\left|1+2\cos\left(\frac{2\pi j}{n}\right)\right|\right)^2. \tag{7.2}$$

For $\frac{j}{n} \leq \frac{1}{8}$, $\cos\left(\frac{2\pi j}{n}\right) \geq \cos\left(\frac{\pi}{4}\right) = \frac{\sqrt{2}}{2}$. Additionally, $2\log(1+\sqrt{2}) > 1$, so by only looking at the first quarter of the terms of the sum, and supposing $p^k \geq 16$ so that these terms exist,

$$\|\operatorname{Log}(u)\|^2 \geq \left\lfloor \frac{p^k-1}{8} \right\rfloor (2\log(1+\sqrt{2}))^2 \geq \left\lfloor \frac{p^k-1}{8} \right\rfloor \geq \frac{p^k}{16}.$$

For $p^k < 16$ we can easily check that the lower bound is still true.

For the upper bound, simply note that for any real x, $\log|1+2\cos(x)| \leq \log 3$, from which we get

$$\|\operatorname{Log}(u)\|^2 \leq \sum_{j\in E}(2\log 3)^2 = 4\#E\log^2 3 = 2\varphi(n)\log^2 3 < 2n\log^2 3,$$

taking the square root we conclude. □

All in all, we get a vector of Λ of order of magnitude $\Theta(\sqrt{n})$, demonstration a large gap between the behaviour of the log-unit lattice and the predictions from the Gaussian heuristic.

Theorem 7.2.5. *Let $n = p^k$ be a prime power with $p \neq 3$. If Λ denotes the log-unit lattice of the n-th cyclotomic number field, we have*

$$\frac{\lambda_1(\Lambda)}{\lambda} = O(n^{-\frac{1}{2}}),$$

where λ denotes the shortest vector length predicted by the Gaussian heuristic.

Proof. From Lemma 7.2.3 we get that the vector $\operatorname{Log}(u)$ is in Λ, for $u = 1+\zeta_n+\zeta_n^{-1}$. From Lemma 7.2.4 we get that $\|\operatorname{Log}(u)\| = O(\sqrt{n})$, therefore $\lambda_1(\Lambda) = O(\sqrt{n})$. We have already discussed why $\lambda = \Theta(n)$, therefore we get the desired result. □

Remark 7.2.6. Our proof is valid for all prime-power cyclotomics with $p \neq 3$. If $n = 3^k$, we can derive similar bounds on the size of the log-unit $\operatorname{Log}(u) = \operatorname{Log}(1+\zeta_n)$.

Figure 7.1 has two purposes, first it illustrates Theorem 7.2.1 and Equation 7.1 by plotting in green the theoretical value of the predicted λ in $\sqrt{\frac{\varphi(n)n}{4\pi e}}$ and in blue the actual value of $\sqrt{\frac{\varphi(n)/2-1}{2\pi e}}\operatorname{vol}(\Lambda)^{\frac{1}{\varphi(n)/2-1}}$, for prime conductors less than 150. Its second purpose is to confirm Lemma 7.2.4 by plotting in red the true length of $\operatorname{Log}(1+\zeta_n+\zeta_n^{-1})$ for the first primes n. Even though the conductors are still very small and the result of Theorem 7.2.5 is asymptotic, we can already see that predictions in blue behave linearly, whereas some shorter vectors behave as $\sqrt{n}$. Experiments were conducted with SageMath [The22].

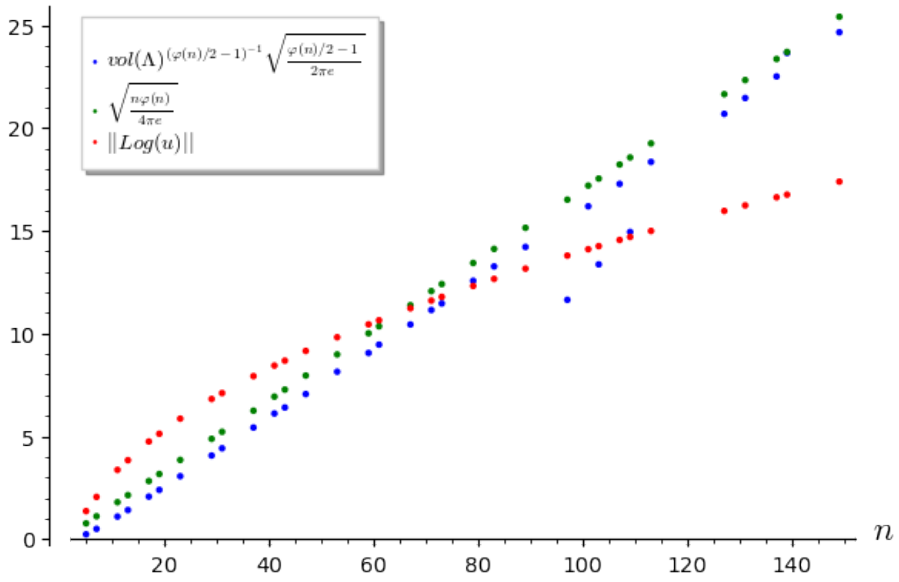

Figure 7.1: Comparing $\|\operatorname{Log}(u)\|$ with the Gaussian heuristic prediction for $\lambda_1(\Lambda)$

7.3 A Bold Conjecture?

As we have seen since the start of this chapter, [BL21] argues that the log-S-unit lattice does not behave like a typical random lattice, meaning that the analysis of the S-unit attack in [PMHS19] does not encompass the whole truth, as the heuristic assumptions needed for the Laarhoven algorithm from [Laa17] do not hold in the same way. In fact, the lattice is more orthogonal, therefore it should reduce much better. The following **controversial** result was conjectured by Daniel J. Bernstein in a 2021 talk at SIAM [Ber21] presenting joint work with Eisenträger, Lange, Rubin, Silverberg and van Vredendaal on S-unit attacks.

Conjecture 7.3.1 ([Ber21])**.** *Under GRH and classical number-theoretic assumptions on the distribution of smooth-norm ideals in power-of-two cyclotomic fields, there exists a quantum algorithm that runs in time* $\exp(\tilde{O}(\sqrt{n}))$ *and solves Ideal-*SVP *with approximation factor* $\gamma = \tilde{O}(\text{poly}(n))$.

Remark 7.3.2. It is not clear if the conjecture should also be extended to other prime-power cyclotomics, namely smooth-degree prime-power cyclotomics, i.e. prime-power cyclotomics for small primes. However as most cryptosystems rely on the power-of-two case, the conjecture as stated above could already have some serious consequences.

Recall that so far (see Table 5.3), the best additive or multiplicative algorithms that find a polynomial approximation factor Ideal-SVP run in exponential time. If

Conjecture 7.3.1 is true, Ideal-γ-SVP with a polynomial approximation factor can be solved in subexponential time for power-of-two cyclotomics. This renders the Ring-LWE to Ideal-SVP security reduction from Theorem 4.3.4 useless: Ring-LWE is at least as hard as Ideal-SVP, however Ideal-SVP is not difficult. Before we panic, it is important to note a few things:

- Conjecture 7.3.1 is my own interpretation of the conjectured result, and may inaccurately reflect the exact assumptions made by its authors, most likely by being slightly weaker.
- Conjecture 7.3.1 has never before been formally stated, analysed or properly justified.
- An attack on Ideal-SVP is not known to directly transpose into an attack on Ring-LWE, therefore the attack would only undermine the power of its security reduction.
- NIST Lattice-based standardised schemes rely on Module-LWE, which is at least as hard as Ring-LWE, and would not be directly endangered by the attack.

After the initial talk [Ber21], the conjecture was very recently mentioned again at the end of an invited talk by Lange at ANTS in August 2022 [Lan22], and briefly in another paper by Bernstein at the same conference [Ber22a]. This last paper is accompanied by software testing a simplified version of the claimed attack, on prime-cyclotomic number fields of degree $23 \leq p \leq 43$. Unfortunately, these experiments are conducted on fields of too small degrees to conclude anything meaningful about the asymptotic claim. The software and its explanation can be found at [Ber22b]. In this final part of the dissertation, we attempt to reproduce the explanations from [Ber21], [Ber22a] and [Ber22b] that could corroborate Conjecture 7.3.1.

The proposed algorithm is for the n-th cyclotomic number field, where n is smooth, in the sense that it has only small prime factors. This is in particular true for powers of two, which is important in the cryptographic setting. The key difference with previous S-unit attacks such as [PMHS19] is to allow for a larger choice of prime ideals in S. Fix $S = \{\mathfrak{p} \text{ prime} \mid \mathcal{N}(\mathfrak{p}) \leq y\}$ where y is a smoothness bound of size $\exp(n^{1/2+o(1)})$. The first step is to precompute small S-units of $\mathcal{O}_K$, i.e. those whose Log-norm squared is less than $n^{1/2+o(1)}$.

Heuristic 7.3.3. *In a smooth-degree cyclotomic number field of conductor n, let $y = \exp(n^{1/2+o(1)})$ and $S = \{\mathfrak{p}$ prime $\mid \mathcal{N}(\mathfrak{p}) \leq y\}$. Then the number of S-units of* Log*-norm squared less than $n^{1/2+o(1)}$ is $\exp(n^{1/2+o(1)})$.*

Heuristic 7.3.3 predicts that there should be approximately $\exp(n^{1/2+o(1)})$ small S-units. The process of finding them is a sieving process that Bernstein calls *filtering*: taking small ring elements and passing them into a filter that checks if they are S-units or not. His very recent paper [Ber22a] facilitates the filtering process, by giving an algorithm that computes the algebraic norm of an element of $\mathcal{O}_K$ of Log-norm squared less than $n^{1/2+o(1)}$ in time $\tilde{O}(n)$, in the case of smooth-degree number fields.

Heuristic 7.3.4. *In a smooth-degree cyclotomic number field of conductor n, given $\exp(n^{1/2+o(1)})$ random S-units of* Log*-norm squared less than $n^{1/2+o(1)}$, they span the full S-unit group with almost certain probability.*

With a subexponential y in an NFS-like way, heuristic 7.3.4 says that the precomputation of a database of $y^{(1+o(1))}$ S-units is enough to get generators of the S-unit group. Bernstein notes that many speed-up tricks can be added to the mix in order to compute small-norm S-units faster, ie by explicitely constructing them with Jacobi sums, or deriving more S-units from those known already by exploiting subfield structure and using automorphisms. This concludes the pre-processing phase.

Given an ideal $\mathfrak{a}$ whose short vectors we would like to find, the online phase of the algorithm proposed by Bernstein uses a single call to the usual quantum algorithm of [BS16] to compute a generator g of a principal ideal equal to $\mathfrak{a}$ times powers of S-units.

It then replaces g by gu/v where u, v are from the precomputed database and are such that $\mathrm{Log}(g) + \mathrm{Log}(u) - \mathrm{Log}(v)$ is close to $\mathrm{Log}(\mathfrak{a})$, makes sure that $g \in \mathfrak{a}$ by multiplying by the appropriate prime ideals and their conjugates, and repeats this process y times. This is very similar to the online phase of Laarhoven's CVP with preprocessing algorithm from [Laa17]. Finally, the algorithm outputs g, which should be a short element of $\mathfrak{a}$.

Heuristic 7.3.5. *In a smooth-degree cyclotomic number field of conductor n, given $\exp(n^{1/2+o(1)})$ S-units of* Log*-norm squared less than $n^{1/2+o(1)}$ that span the full S-unit, an S-generator g of an ideal $\mathfrak{a}$ of $\mathcal{O}_K$ reduces via the log-S-unit lattice in time $\exp(\tilde{O}(\sqrt{n}))$ to an element of $\mathfrak{a}$ of* Log*-norm within a polynomial factor of the shortest length $\lambda_1(\mathfrak{a})$.*

Heuristic would ensure the validity of the main conjecture. As things stand, it seems that Heuristic 7.3.3 is very likely and is supported by various experiments, and that Heuristic 7.3.4 is also likely as experimentally and under GRH, problems seem to arise only when the number fields are unbalanced (e.g. not in the smooth case). However the controversial point is the analysis of the reduction in Heuristic 7.3, where all attempts to analyse rely on lattice heuristics that are very difficult to justify properly.

Chapter 8

Conclusion

In this dissertation, we have surveyed the state-of-the-art classical and quantum algorithms that solve the Approximate Shortest Vector Problem in ideal lattices. We have seen that the hardness of this problem is at the foundation of the security of lattice-based post-quantum cryptosystems that are to be standardised. General lattices are very well studied and benefit from strong security guarantees, but recent works have shown that because the lattices used in modern cryptosystems have a lot more structure than general lattices, stronger algorithms exploiting this structure and relying on quantum computers exist.

These so-called unit or S-unit attacks rely on special lattices, the log-unit and log-S-unit lattices. Previous works analysed the speed of S-unit attacks using the assumption that these lattices could be studied as if they were generated randomly. We extended recent results that quantify how different the log-unit lattice really is from the random models, in the case of prime-power cyclotomic fields. These results could foreshadow that the power of S-unit attacks has been underestimated, and that the security of supposedly quantum-resistant schemes that are being deployed today is not as strong as we thought.

However, the conjectured subexponential attack on Ideal-SVP with polynomial approximation factor remains very poorly documented and is believed by some experts to be exaggerated. Nevertheless, it makes for a burning topic of debate in the post-quantum cryptography community. The field is evolving fast, and will have massive consequences on global privacy and information security in the future. It lies at the intersection between many areas in science, is still at a very early stage, and would benefit enormously from greater and more coordinated efforts towards cryptanalysis: lots of aspects of structured lattices and S-unit attacks remain to be understood, and for that researchers with backgrounds in number theory, quantum computing and cryptography will all have to work towards the same objectives.

Bibliography

[AD17] M.R. Albrecht and A. Deo. Large modulus ring-lwe $\geq$ module-lwe. Cryptology ePrint Archive, Paper 2017/612, 2017. https://eprint.iacr.org/2017/612.

[Ajt96] M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96, page 99–108, New York, NY, USA, 1996. Association for Computing Machinery.

[Ajt98] M. Ajtai. The shortest vector problem in l2 is np-hard for randomized reductions (extended abstract). In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, STOC '98, page 10–19, New York, NY, USA, 1998. Association for Computing Machinery.

[AM07] H. Abbaspour and M.A. Moskowitz. *Basic lie theory*. World Scientific, Hackensack, N.J., 2007. OCLC: ocn141187986.

[Bab86] L. Babai. On Lovász' lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, March 1986.

[Bac90] E. Bach. Explicit bounds for primality testing and related problems. *Mathematics of Computation*, 55(191):355–380, 1990.

[BEF+17] J.-F. Biasse, T. Espitau, P.-A. Fouque, A. Gélin, and P. Kirchner. Computing generator in cyclotomic integer rings. In *Advances in Cryptology-EUROCRYPT 2017*, volume 10210 of *Lecture Notes in Computer Science*, pages 60–88, Paris, France, April 2017.

[Ber14] D.J. Bernstein. Blogpost: "a subfield-logarithm attack against ideal lattices", 2014. http://blog.cr.yp.to/20140213-ideal.html, Last updated 2022-01-09.

[Ber21] D.J. Bernstein. Talk: "s-unit attacks", 2021. at SIAM Annual Meeting 2021, https://cr.yp.to/talks.html#2021.08.20.

[Ber22a] D.J. Bernstein. Fast norm computation in smooth-degree abelian number fields. Cryptology ePrint Archive, Paper 2022/980, 2022. https://eprint.iacr.org/2022/980.

[Ber22b] D.J. Bernstein. Website: "s-unit attacks: filtered", 2022. https://s-unit.attacks.cr.yp.to/filtered.html, Last updated 2022-07-31.

[BF14] J.-F. Biasse and C. Fieker. Subexponential class group and unit group computation in large degree number fields. *LMS Journal of Computation and Mathematics*, 17:385–403, 1 2014.

[BGV11] Z. Brakerski, C. Gentry, and V. Vaikuntanathan. Fully homomorphic encryption without bootstrapping. Cryptology ePrint Archive, Paper 2011/277, 2011. https://eprint.iacr.org/2011/277.

[BL21] D.J. Bernstein and T. Lange. Non-randomness of s-unit lattices. Cryptology ePrint Archive, Report 2021/1428, 2021. https://ia.cr/2021/1428.

[BLNRL21] O. Bernard, A. Lesavourey, T.-H. Nguyen, and A. Roux-Langlois. Log-s-unit lattices using explicit stickelberger generators to solve approx ideal-svp. Cryptology ePrint Archive, Paper 2021/1384, 2021. https://eprint.iacr.org/2021/1384.

[BPR] J. Buhler, C. Pomerance, and L. Robertson. Heuristics for class numbers of prime-power real cyclotomic fields.

[BRL20] O. Bernard and A. Roux-Langlois. Twisted-phs: Using the product formula to solve approx-svp in ideal lattices. In S. Moriai and H. Wang, editors, *Advances in Cryptology – ASIACRYPT 2020*, pages 349–380, Cham, 2020. Springer International Publishing.

[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 27th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 893–902. SIAM, January 2016.

[CD22] W. Castryck and T. Decru. An efficient key recovery attack on sidh (preliminary version). Cryptology ePrint Archive, Paper 2022/975, 2022. https://eprint.iacr.org/2022/975.

[CDPR16] R. Cramer, L. Ducas, C. Peikert, and O. Regev. Recovering short generators of principal ideals in cyclotomic rings. In M. Fischlin and J.-S. Coron, editors, *Advances in Cryptology – EUROCRYPT 2016*, pages 559–585, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.

[CDW17] R. Cramer, L. Ducas, and B. Wesolowski. Short stickelberger class relations and application to ideal-svp. In J.-S. Coron and J.B. Nielsen, editors, *Advances in Cryptology – EUROCRYPT 2017*, pages 324–348, Cham, 2017. Springer International Publishing.

[CDW21] R. Cramer, L. Ducas, and B. Wesolowski. Mildly short vectors in cyclotomic ideal lattices in quantum polynomial time. *J. ACM*, 68(2), jan 2021.

[CGS14] P. Campbell, M. Groves, and D. Shepherd. Soliloquy: A cautionary tale. ETSI 2nd Quantum-Safe Crypto Workshop, 2014. Available at http://docbox.etsi.org/Workshop/2014/201410_CRYPTO/S07_Systems_and_Attacks/S07_Groves_Annex.pdf.

[CN11] Y. Chen and P.Q. Nguyen. Bkz 2.0: Better lattice security estimates. In *ASIACRYPT*, 2011.

[DPW19] L. Ducas, M. Plançon, and B. Wesolowski. On the shortness of vectors to be found by the ideal-svp quantum algorithm. Cryptology ePrint Archive, Report 2019/234, 2019. https://ia.cr/2019/234.

[EHKS14] 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 Forty-Sixth Annual ACM Symposium on Theory of Computing*, STOC '14, page 293–302, New York, NY, USA, 2014. Association for Computing Machinery.

[GM03] D. Goldstein and A. Mayer. On the equidistribution of hecke points. *Forum Mathematicum*, 15:165–189, 2003.

[GMSS99] O. Goldreich, D. Micciancio, S. Safra, and J.-P. Seifert. Approximating shortest lattice vectors is not harder than approximating closest lattice vectors. *Information Processing Letters*, 71(2):55–61, 1999.

[HPS98] J. Hoffstein, J. Pipher, and J.H. Silverman. Ntru: A ring-based public key cryptosystem. In J.P. Buhler, editor, *Algorithmic Number Theory*, pages 267–288, Berlin, Heidelberg, 1998. Springer Berlin Heidelberg.

[Koc00] H. Koch. *Number theory: algebraic numbers and functions*. Number v. 24 in Graduate studies in mathematics. American Mathematical Society, Providence, RI, 2000.

[Laa17] T. Laarhoven. Sieving for closest lattice vectors (with preprocessing). In *Lecture Notes in Computer Science*, pages 523–542. Springer, 2017.

[Lan85] S. Lang. *SL2(R)*. Graduate texts in mathematics. Springer, 1985.

[Lan22] T. Lange. Talk: "s-unit attacks", 2022. at ANTS-XV,https://www.youtube.com/watch?v=AJB-_fYAJmE&ab_channel=ANTS.

[LLL82] H.W. jr. Lenstra, A.K. Lenstra, and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261:515–534, 1982.

[LPR10] V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *Advances in Cryptology – EUROCRYPT 2010*, pages 1–23, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.

[LS12] A. Langlois and D. Stéhle. Worst-case to average-case reductions for module lattices. Cryptology ePrint Archive, Paper 2012/090, 2012. https://eprint.iacr.org/2012/090.

[Mic01] D. Micciancio. The shortest vector problem is NP-hard to approximate to within some constant. *SIAM Journal on Computing*, 30(6):2008–2035, March 2001. Preliminary version in FOCS 1998.

[Mic07] D. Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Comput. Complex.*, 16(4):365–411, dec 2007.

[Ngu10] P. Q. Nguyen. *Hermite's Constant and Lattice Algorithms*, pages 19–69. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.

[NV09] P. Q. Nguyen and B. Valle. *The LLL Algorithm: Survey and Applications*. Springer Publishing Company, Incorporated, 1st edition, 2009.

[PMHS19] A. Pellet-Mary, G. Hanrot, and D. Stehlé. Approx-svp in ideal lattices with pre-processing. Cryptology ePrint Archive, Report 2019/215, 2019. https://ia.cr/2019/215.

[Reg05] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing - STOC '05*, page 84, Baltimore, MD, USA, 2005. ACM Press.

[Rog56] C.A. Rogers. The Number of Lattice Points in a Set. *Proceedings of the London Mathematical Society*, s3-6(2):305–320, April 1956.

[Sch87] C.P. Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theoretical Computer Science*, 53(2-3):201–224, 1987.

[Sie45] C.L. Siegel. A Mean Value Theorem in Geometry of Numbers. *The Annals of Mathematics*, 46(2):340, April 1945.

[SS16] A. Strömbergsson and A. Södergren. On the generalized circle problem for a random lattice in large dimension, 2016.

[SSTX09] D. Stehlé, R. Steinfeld, K. Tanaka, and K. Xagawa. Efficient Public-Key Encryption Based on Ideal Lattices (Extended Abstract). In *Asiacrypt 2009*, pages 617–635, Japan, 2009.

[ST02] I. Stewart and D. Tall. *Algebraic number theory and Fermat's last theorem*. AK Peters, Natick, Mass, 3rd ed edition, 2002.

[SV09] N.P. Smart and F. Vercauteren. Fully homomorphic encryption with relatively small key and ciphertext sizes. Cryptology ePrint Archive, Paper 2009/571, 2009. https://eprint.iacr.org/2009/571.

[The22] The Sage Developers. *SageMath, the Sage Mathematics Software System (Version 9.2)*, 2022. https://www.sagemath.org.

[Was97] L. C. Washington. *Introduction to Cyclotomic Fields*, volume 83 of *Graduate Texts in Mathematics*. Springer New York, New York, NY, 1997.