NTNU

Norwegian University of
Science and Technology

Hardness of Lattice Based Cryptography

Thor Tunge

Master of Science in Mathematical Sciences
Submission date: June 2018
Supervisor: Kristian Gjøsteen, IMF

Norwegian University of Science and Technology
Department of Mathematical Sciences

Hardness of Lattice Based Cryptography

Thor Tunge

June 1, 2018

Abstract

In this thesis we will discuss hard computational problems in lattice theory and relate them to cryptographic constructions. Many of these problems enjoy average-case hardness which makes them attractive for cryptography. In addition, lattice based cryptography is a candidate for post-quantum cryptography, as there is no known quantum algorithm which breaks various hardness theorems.

We build a foundation in algebraic number theory to have the required background to discuss schemes based on discrete algebraic structures. These structures are free $\mathbb{Z}$-modules which permits unique factorization in prime ideals. We relate this algebraic number theory to lattices in $\mathbb{R}^n$ so we can use the theory from algebra to our advantage.

We then define some standard hard computational lattice problems and show how many of these are related to each other. We prove that these problems are at least as hard as finding the shortest vector of a lattice, which we conjecture is computationally infeasible. We then prove a quantum reduction the learning with errors problem, a problem in machine learning . We also show that there is a similar reduction for a variant of this problem over more general rings.

Sammendrag

I denne oppgaven drøfter vi vanskelige problemer i lattice teori og kobler dem opp mot kryptografiske konstruksjoner. Mange av disse problemene er like vanskelige i 'average-case' som i 'worst-case'. I tillegg er lattice-basert kryptografi en potensiell mulighet som sikker cryptografi mot kvante-datamaskiner siden vi ikke har noen algoritme som bryter teoremer for sikkerhet.

Vi begynner med å bygge opp teori fra algebraisk tallteori for å ha den nødvendige bakgrunnen til å diskutere kryptografi basert på diskrete algebraiske konstruksjoner. Disse konstruksjonene er frie $\mathbb{Z}$-moduler som har unik faktorisering i prim-idealer. Vi relaterer disse til latticer i $\mathbb{R}^n$.

Deretter definerer vi standard, vanskelige lattice problemer og viser hvordan de relaterer til hverandre. Vi viser at disse problemene er mist like vanskelige som å finne den korteste vektoren i en lattce, noe som vi anntar er vanskelig. Deretter beviser vi en kvante-reduksjon fra et maskin-læring problem kalt 'learning with errors'. Vi viser også en tilsvarende reduksjon fra en variant over mer generelle ringer.

Acknowledgments

I would like to thank my advisor Sverre Olaf Smalø for helping me with building a good theoretical foundation for my thesis and always keeping his door open. I would also like to thank my co-advisor Kristian Gjøsteen for providing feedback whenever I needed and allowing me to pursue the aspect of cryptography I wanted. I would like to thank Martin Strand for additional feedback and discussion related to cryptography.

Contents

1 Introduction

A lattice is a discrete additive subgroup of $\mathbb{R}^n$. By viewing the basis-vectors for a basis of $\mathbb{R}^n$ as column vectors in a matrix B, a lattice Λ generated by B is any linear combination of the columns of B with integer coefficients.

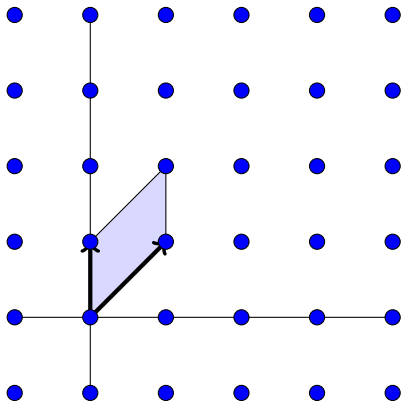

Figure 1: Lattice with basis $B = \{(1,1),(0,1)\}$. This lattice is isomorphic to the lattice with basis vectors $(1,0)$ and $(-1,1)$ (and countably many others). In some sense, this is the most 'basic' lattice because it is equal to $\mathbb{Z}^2$.

Consider a simple encryption scheme where a lattice point $\boldsymbol{x} = B\boldsymbol{z}$ is given and we encrypt it by permuting it by a real vector $\boldsymbol{e}$

$$\text{Encrypt}(\boldsymbol{x}) = \boldsymbol{x} + \boldsymbol{e} \notin \Lambda.$$

To decrypt such a message we need to find the closest lattice vector to $\boldsymbol{x}+\boldsymbol{e}$, that is, solve the *closest vector problem* on Λ. Closely related to this is the *shortest vector problem* which is to find the shortest vector in a given lattice. How hard these problems are depends on how we choose the basis B for Λ. By providing a 'good' basis as the private key and 'bad' basis for the public key we can make ciphertexts easy to decrypt but hard to attack. We see that for this to work we need that the problem of finding the closest lattice vector to an arbitrary point in $\mathbb{R}^n$ must be hard. This is one of the problems we define in Section 4 and show that is at least as hard as conjectured hard problems.

There is a problem in machine learning called *learning with errors* (LWE).

In short, it is to recover a secret $\boldsymbol{s} \in \mathbb{Z}_q^n$ given equations

$$\begin{aligned} b_1 &= \langle \boldsymbol{a}_1, \boldsymbol{s} \rangle + e \mod q \\ b_2 &= \langle \boldsymbol{a}_2, \boldsymbol{s} \rangle + e \mod q \\ &\vdots \end{aligned}$$

for $\boldsymbol{a}_i \xleftarrow{r} \mathbb{Z}_q^n$ and errors $e \leftarrow \chi$ for a distribution χ over $\mathbb{R}$. A simple scheme based on this problem can be shown to be semantically secure. In addition, we show a quantum reduction from hard lattice problems to LWE which strengthens our confidence in schemes based on LWE.

The hardness of these lattice problems lay the foundations for many cryptographic schemes. Most notably is the NTRU (N-Th Degree Truncated Polynomial Ring) scheme which is known to be insecure if an attacker can find sufficiently short lattice vectors[HPS98]. In 2005, Oded Regev introduced a simple scheme whose security is based on the LWE problem, and showed it to be to be semantically secure if LWE is hard. He also proved a quantum polynomial reduction from standard lattice problems to LWE, strengthening the claim that the scheme is secure. This was later generalized to other rings, and the corresponding problems enjoy similar hardness theorems[LPR10].

In 2009 Craig Gentry introduced the first so-called fully homomorphic scheme, a method of encrypting such that both addition and multiplication is a homomorphic operation under the encryption map[Gen09], and proved that it could perform arbitrarily many operations without decryption errors. His ideas has been used to create other fully homomorphic schemes[BGV14, vDGHV10, LATV17]. Many of these schemes are based on the hardness of lattice problems and the LWE problem. A good understanding of the security of lattice based security therefore translates to provably secure fully homomorphic encryption schemes.

While cryptographic primitives based on integer factorization and discrete logarithms have been shown to be insecure against quantum computers[Sho97], there have been little progress in attacking lattice based cryptography with quantum computers. This makes lattice based cryptography an attractive alternative in a post-quantum world.

1.1 Overview

In Section 2 we introduce the basic algebraic number theory required for this thesis. We define discrete algebraic structures which permits unique factorization and some properties of important maps. Next, we define lattices in $\mathbb{R}^n$ and relate them to the algebraic number theory by showing that certain algebraic structures can be seen as lattices in $\mathbb{R}^n$. In Section 4 we define standard lattice problems, such as the *shortest vector problem* (SVP) and *closest vector problem* (CVP). For other lattice problems we show how they relate to these conjectured hard problems. In Section 5 we define the *learning with errors* (LWE) problem and show a quantum reduction from SVP to LWE.

When discussing various lattice problems we do it in the following way: First we define the problem, then show how we can attempt to break the problem (subsection *attack*) and lastly how this problem relates to other hard problems (subsection *hardness*). In other words, the attack section describes an algorithm to solve a given problem *without* any oracles while the hardness section describes how we can attack this problem with an oracle.

1.2 Notation

We denote by $||x||$ the Euclidean norm of x, and for simplicity we only use this norm in this thesis. $\lfloor a \rceil$ denotes the closest integer to a, mapping $1/2$ to 1. If $\boldsymbol{a}$ is a vector $\lfloor \boldsymbol{a} \rceil$ denotes this rounding on each coordinate. The notation $q = \text{poly}(n)$ means that q is polynomial in n. A *negligible* function $f : \mathbb{N} \to \mathbb{R}$ is such that $\lim_{n\to\infty} n^c f(n) = 0$ for any $c > 0$, i.e. f is asymptotically smaller than any polynomial.

We will, with slight abuse of notation, denote by $a \xleftarrow{r} A$ to mean that a is sampled uniformly from a set A. If B is a distribution, $e \leftarrow B$ means that e is sampled according to B. To confuse, we will use the standard notation from computer science $t \leftarrow t_1$ to denote that t is assigned the value of t_1. This is done only in algorithms and should be clear from context. An algorithm/reduction A that uses an oracle O is denoted A^O.

2 Algebraic Number Theory

Algebraic number theory is the study of finite extensions of K of $\mathbb{Q}$. We generalize the notion of integers by considering the integral closure of $\mathbb{Z}$ in K, study the ideals in this 'new' ring of integers and extend the notion of ideal to give the set of ideals a group structure. We end up with that the ring of integers $\mathcal{O}$ (and its ideals) permits a unique factorization of prime ideals. Throughout this section (and the thesis in general) K denotes a finite extension of $\mathbb{Q}$ which we will call a *number field*.

In this section we will build the theory required to connect lattices in $\mathbb{R}^n$ to discrete algebraic structures in number fields. In particular, we show that an ideal $\mathcal{I}$ in the ring of integers is a free $\mathbb{Z}$-module and that it can be embedded into $\mathbb{R}^n$ as a lattice.

2.1 Norm, Trace and Geometry

We start by defining the norm and trace of the elements of K.

Definition 2.1 (Norm and Trace)**.** Let $K/\mathbb{Q}$ be a finite field extension. Consider the multiplication map

$$\begin{aligned} \mu_\alpha : & K \to K \\ & x \mapsto \alpha x. \end{aligned}$$

By fixing a $\mathbb{Q}$-basis of K, we can define the *norm* and *trace* maps to be

$$\begin{aligned} \mathrm{N}_{K/\mathbb{Q}}(\alpha) &= \det(\mu_\alpha) \\ \mathrm{Tr}_{K/\mathbb{Q}}(\alpha) &= \mathrm{Tr}(\mu_\alpha), \end{aligned}$$

respectively. It is easy to verify that the norm map is multiplicative while the trace map is additive. Both maps are invariant under choice of basis. Notice that is makes sense to fix a $\mathbb{Q}$ basis since the extension is finite.

A number field $K = \mathbb{Q}(\zeta)$ of index n has exactly n ring embeddings $\sigma_i : K \to \mathbb{C}$, by sending ζ to a root of the minimal polynomial $f(X) \in \mathbb{Q}[X]$ of ζ. Since the complex roots come in conjugate pairs, so does the embeddings. Denote by s_1 the number of real embeddings and by s_2 the number of *pairs* or complex embeddings such that $n = s_1 + 2s_2$. By convention, order the embeddings as follows: $\sigma_1, \dots \sigma_{s_1}$ are the real embeddings. The complex ones is ordered such that $\sigma_{s_1+s_2+j} = \overline{\sigma_{s_1+j}}$ for $j \in \{1, \dots, s_2\}$. Therefore, with

increasing j we have first the s_1 real embeddings, then the s_2 complex embeddings *with no conjugates among themselves* and lastly the s_2 conjugates. Now we define

$$\sigma(x) = (\sigma_1(x), \ldots, \sigma_n(x)) \in \mathbb{C}^n$$

as the *canonical* embedding of $K \mapsto \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$. We can now get some geometry on K.

Definition 2.2. The length of an element $x \in K$ is given by

$$||x|| := ||\sigma(x)||$$

where σ is the canonical embedding.

Notice that this means that certain elements of K will have unusual norms, e.g. let $K = \mathbb{Q}(\zeta_p)$ be a cyclotomic field. A root of unity ζ_p, which has Euclidean norm 1 when viewed as an element in $\mathbb{C}$, will be embedded as

$$\sigma(\zeta_p) = (\zeta_p, \zeta_p^2, \ldots, \zeta_p^{n-1})$$

which means that its norm is $||\zeta_p|| = ||(\zeta_p, \ldots, \zeta_p^{n-1}||_2 = \sqrt{n}$. It can be shown that the norm and trace of an element in K is given by embeddings,

$$\mathrm{Tr}(x) = \sum_{i=1}^{n} \sigma_i(x)$$
$$\mathrm{N}(x) = \prod_{i=1}^{n} \sigma_i(x).$$

It is not hard to see that this gives us

$$\mathrm{Tr}(x \cdot y) = \sum_{i=1}^{n} \sigma_i(x)\sigma_i(y) = \langle \sigma(x), \overline{\sigma(y)} \rangle,$$

which means that the trace map is a bilinear form like the inner product on $\mathbb{C}^n$.

2.2 Ring of Integers

The ring of integers is the integral closure of $\mathbb{Z}$ in K, i.e. any roots of monic polynomials with integral coefficients. To avoid ambiguity we call $\mathbb{Z}$ the *rational* integers and its integral closure *algebraic* integers. We drop these prefixes when the ring considered is clear from context.

Definition 2.3 (Integral Element and Ring of Algebraic Integers)**.** Let K be a finite extension field of $\mathbb{Q}$. We say that α is an *integral element* in K if α is a root of a monic polynomial with rational integer coefficients. The *ring of algebraic integers* of K, denoted $\mathcal{O}$ is the set of all integral elements of K.

It is not obvious that $\mathcal{O}$ is a ring, but the we will prove this foreshadowing later. Trivially, a rational integer m is the root of $x - m$ and therefore we have $\mathbb{Z} \subseteq \mathcal{O}$. We also have that the ring of integers of $K = \mathbb{Q}$ is simply $\mathbb{Z}$, which is why we call $\mathbb{Z}$ the rational integers. We note that the ring of integers is often denoted $\mathcal{O}_K$ because it depends on the extension field. This will not be important for our purposes as the field K is always clear from context, so we omit this subscript.

It can be shown that the norm and trace of any element $\alpha \in \mathcal{O}$ is an integer, we omit the proof.

Proposition 2.4. The norm and trace of $\alpha \in \mathcal{O}$ are rational integers.

Additionally, $N(\alpha) = \pm 1 \Leftrightarrow \alpha \in \mathcal{O}^*$: The norm of an element is ± 1 if, and only if, it is a unit in the ring of integers. We can try to prove that $\mathcal{O}$ is a ring in the following way: Given two elements algebraic integers α_1, α_2 we can try to find monic polynomials f, g with integer coefficient explicitly such that $f(\alpha_1\alpha_2) = 0$ and $g(\alpha_1 + \alpha_2) = 0$. However, this is not easy, so we prove this is a different way. We devote the rest of this subsection to prove this, in addition to showing that $\mathcal{O}$ is free as a $\mathbb{Z}$-module.

Proposition 2.5. The minimal polynomial of $\alpha \in K$ has rational integer coefficients if, and only if, α is an algebraic integer.

Proof. Assume that the minimal polynomial of α, say $g(X)$ has rational integer coefficients. Since $g(\alpha) = 0$ and g is monic by the definition of a minimal polynomial, α is an algebraic integer.

Assume now that α is an algebraic integer. By definition, there exists an $f(X) \in \mathbb{Z}[X]$ such that $f(\alpha) = 0$. If $f(X)$ is the minimal polynomial

we are done. Let therefore $g(X)$ be the minimal polynomial of α. We need to show that $g(X)$ also has integer coefficients. Since $g(X)$ is the minimal polynomial, we have that

$$f(X) = g(X)h(X)$$

for some $h(X) \in \mathbb{Q}[X]$. If $g(X)$ has a rational coefficient, then one of them must be divisible by p (by fundamental theorem of algebra). Assume that $g(X)$ has rational coefficients, i.e. that at least one denominator is divisible by a prime p. Let u be the smallest integer such that $p^u g(X)$ has no denominators divisible by p. Similarly, let v be the smallest integer such that $p^v h(X)$ has no denominators divisible by p. Now, the left side of

$$p^u g(X) p^v h(X) = p^{u+v} f(X) \tag{1}$$

has no denominators divisible by p. If we regard (1) as an equation in $\mathbb{Z}_p[X]$, since $f(X)$ has integer coefficients, the right side is 0. Since we have removed all $p = 0 \in \mathbb{Z}_p$ from the denominators of left side, regarding it as polynomials in $\mathbb{Z}_p$ makes sense. Therefore,

$$p^u g(X) p^v h(X) = 0 \in \mathbb{F}_p[X],$$

and because we chose u and v to be minimal, neither $p^u g(X)$ nor $p^v h(X)$ are zero-polynomials. Since $\mathbb{F}_p[X]$ has no zero divisors, this leads to a contradiction and we conclude that $g(X) \in \mathbb{Z}[X]$. □

Towards a goal of proving that $\mathcal{O}$ is a ring, we use a relationship between an algebraic integer α and $\mathbb{Z}[\alpha]$.

Proposition 2.6. Let $\alpha \in K$. Then α is an algebraic integer if, and only if, $\mathbb{Z}[\alpha]$ is finitely generated as a $\mathbb{Z}$-module.

Proof. Assume α is an algebraic integer, and let $g(X)$ be its minimal polynomial. Let $\deg(g) = m$. By Proposition 2.5 we have that $g(X)$ is monic with integer coefficients and can therefore write

$$g(X) = X^m + \hat{g}(X)$$

for some $\hat{g}(X)$. Because $g(\alpha) = 0$ we can write $\alpha^m = -\hat{g}(\alpha)$ where $\deg \hat{g}(X) < \deg g(X)$. This means that any α^u can be written as a $\mathbb{Z}$-linear combination of $\{1, \alpha, \alpha^2, \ldots, \alpha^{m-1}\}$ which therefore generate $\mathbb{Z}[\alpha]$, i.e. $\mathbb{Z}[\alpha]$ is free.

Now assume that $\mathbb{Z}[\alpha]$ is free as a $\mathbb{Z}$-module with basis elements $\{a_0, a_1, \ldots, a_{m-1}\}$. Let $f_i(X), i = 0, \ldots m-1$ be such that $a_i = f(\alpha)$ for an $\alpha \in K$. Now pick an $N > \deg f_i$ for all $i = 0, \ldots, m-1$. Since $\mathbb{Z}[\alpha]$ is free

$$\alpha^N = \sum_{i=0}^{m-1} a_i b_i \quad \text{for some } b_i \in \mathbb{Z}.$$

Choose

$$f(X) = X^N - \sum_{i=1}^{m-1} f_i(X) b_i.$$

Because we chose N to be larger than all $\deg f_i(X)$, $f(X)$ is monic and has integer coefficients. Furthermore, $f(\alpha) = 0$ so α is an algebraic integer. □

We now have the required machinery to prove that the ring of integers is indeed a ring.

Theorem 2.7. *The ring of integers $\mathcal{O}$ of a finite extension field K of $\mathbb{Q}$ is a ring.*

Proof. Let $\alpha, \beta \in \mathcal{O}$. By Proposition 2.6 we have that $\mathbb{Z}[\alpha]$ and $\mathbb{Z}[\beta]$ are finitely generated, and therefore $\mathbb{Z}[\alpha, \beta]$ is finitely generated. Regarding $\mathbb{Z}[\alpha, \beta]$ as a ring, we have that $\alpha\beta, \alpha \pm \beta \in \mathbb{Z}[\alpha, \beta]$. Now, since $\mathbb{Z}[\alpha\beta]$ and $\mathbb{Z}[\alpha \pm \beta]$ are both subrings of $\mathbb{Z}[\alpha, \beta]$, they are finitely generated. By Proposition 2.6 again, we conclude that $\alpha\beta$ and $\alpha \pm \beta$ are algebraic integers and $\mathcal{O}$ is therefore a ring. □

If we let the field be $\mathbb{Q}$ then the ring of integers equals the rational integers $\mathbb{Z}$. To see this, Gauss' lemma gives us that every root of a monic polynomial with rational coefficients is a rational integer. As $\mathbb{Z} \subseteq \mathcal{O}$ this is the simplest form the ring of integers can have. $\mathcal{O}$ might, however, be more complicated than this. Without any restrictions on the field K, the ring of integers can be quite hard to determine. We use the rest of most of this section to determine some of the structure of $\mathcal{O}$ and its ideals.

Proposition 2.8. Let K be a number field with ring of integers $\mathcal{O}$. Then $\mathbb{Q}\mathcal{O} = K$.

Proof. Clearly $\mathbb{Q}\mathcal{O} \subseteq K$.

$K \subseteq \mathbb{Q}\mathcal{O}$: Assume $\alpha \in K$. Let $f(X) \in \mathbb{Z}[X]$ be the minimal polynomial of α. Let d be the least common multiple of the coefficients of $f(X)$. Now, define the polynomial $g(X)$ by

$$g(X) := d^{\deg f(X)} f(\frac{X}{d}).$$

By design, $df(X)$ will have only integer coefficients and $d^{\deg f(X)} f(X/d)$ will be monic. Additionally, $g(\alpha d) = f(\alpha) = 0$ so $g(X)$ is a monic polynomial with integer coefficients that has αd as a root. By Proposition 2.6 we get that $\alpha d \in \mathcal{O} \subseteq \mathbb{Q}\mathcal{O}$ which we wanted to prove. □

From this proof we get a small, but important, corollary:

Proposition 2.9. For any $\alpha \in K$, there exists $d \in \mathbb{Z}$ such that $\alpha d \in \mathcal{O}$.

Proof. See proof of Proposition 2.8. Construct d the same way as in the proof. □

Notice that the $d \in \mathbb{Z}$ acts as a denominator: Multiplying α by d 'cancels the denominators' of the element α. In other words, an element from the extension field $K = \mathbb{Q}(\zeta)$ has 'denominators' which can be canceled by a rational integer to get an algebraic integral element in the ring $\mathcal{O}$. This strengthens the intuition that $\mathcal{O}$ is an extension of the rational integers $\mathbb{Z}$. We have foreshadowed that $\mathcal{O}$ is free and we finally prove it.

Theorem 2.10. *The ring of integers $\mathcal{O}$ is a free abelian group of rank $n = [K : \mathbb{Q}]$*

Proof. We know that there exists a $\mathbb{Q}$-basis $\{\alpha_1, \ldots, \alpha_n\}$ of K with $\alpha_i \in \mathcal{O}$ for all $i = 1, \ldots n$, namely take a $\mathbb{Q}$-basis of K and multiply each basis element with an appropriate d to get an element in $\mathcal{O}$ as per Proposition 2.9. By fixing the basis $\{\alpha_1, \ldots, \alpha_n\}$, there exists a unique basis $\{\hat{\alpha}_1, \ldots, \hat{\alpha}_n\} \subseteq K$ such that $\langle \alpha_i, \hat{\alpha}_j \rangle = \delta_{ij}$ where δ_{ij} is the Kronecker delta. Now write

$$x = \sum_{i=1}^{n} a_i \hat{\alpha}_i \qquad a_i \in \mathbb{Z}$$

and by the linearity of the inner product that

$$\langle \alpha_i, x \rangle = \langle \alpha_i, \sum_{i=1}^{n} a_i \hat{\alpha}_i \rangle = a_i \quad \text{for all } i = 1, \dots n.$$

Now because $\langle \alpha_i, x \rangle \in \mathbb{Z}$ we get that $a_i \in \mathbb{Z}$ for all $i = 1 \dots n$. Therefore

$$\mathbb{Z}[\alpha_1, \dots, \alpha_n] \subseteq \mathcal{O} \subseteq \mathbb{Z}[\hat{\alpha}_1, \dots, \hat{\alpha}_n]$$

where the first inclusion is obvious. Since $\mathcal{O}$ is a subgroup of a free abelian group, $\mathcal{O}$ is itself a free abelian group. It contains at least n linearly independent elements, namely $\{\alpha_1, \dots, \alpha_n\}$ so its rank is at least n. However, its rank cannot exceed $\mathbb{Z}[\hat{\alpha}_1, \dots, \hat{\alpha}_n]$ so we conclude that the rank of $\mathcal{O}$ is n. $\square$

We also have that any ideal $\mathcal{I} \subseteq \mathcal{O}$ is free of rank n by the following diagram

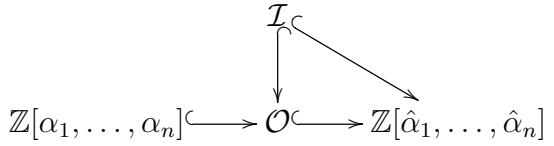

and similar arguments as above. To make this rigorous we need to show that any ideal $\mathcal{I}$ contain at least n linearly independent elements. We omit the details.

2.3 Ideals of Ring of Integers

We now know that $\mathcal{O}$ is a ring which is finitely generated as a $\mathbb{Z}$-module. Note that this also means that $\mathcal{O}$ is Noethrian, which in turn means that any set of ideals of $\mathcal{O}$ contains a maximal element. In this section we want to show that any ideal can be uniquely factored into prime ideals. In addition, we extend the notion of an ideal to create a set of ideal-like object with group structure.

We begin by defining the norm of an ideal.

Definition 2.11 (Ideal Norm)**.** The norm of a non-zero ideal $\mathfrak{a} \subseteq \mathcal{O}$ is

$$N(\mathfrak{a}) = |\mathcal{O}/\mathfrak{a}|$$

It can be shown that the norm of an ideal is finite[Ogg10]. We see that a 'large' ideal in the subset sense will have a small norm, e.g. the 'large' ring $\mathcal{O} \subseteq \mathcal{O}$ has ideal norm 1.

For primes in $\mathbb{Z}$ we have two equivalent definitions: A number is prime if $p = ab \implies a$ or b is a unit *or* $p|ab \implies p|a$ or $p|b$. However, this definition does not generalize[Ogg10]. In the general case, we differentiate between these two properties, and call the first one *irreducible* and the second one *prime*. We define a prime ideal as follows:

Definition 2.12 (Prime Ideal)**.** An ideal $\mathfrak{p}$ is *prime* if, for any elements $a, b \in \mathfrak{p}$,

$$ab \in \mathfrak{p} \Rightarrow a \in \mathfrak{p} \quad \text{or} \quad b \in \mathfrak{p}.$$

A special property of non-zero prime ideals in $\mathcal{O}$ is that they are all maximal.

Proposition 2.13. Every non-zero prime ideal in $\mathcal{O}$ is maximal.

Proof. Any ideal $\mathcal{I} \subseteq \mathcal{O}$ is maximal if, and only if, the quotient $\mathcal{O}/\mathcal{I}$ is a field. We therefore show that $\mathcal{O}/\mathfrak{p}$ is a field. Take $x \in \mathcal{O}/\mathfrak{p}$. Because $\mathfrak{p}$ is prime, $\mathcal{O}/\mathfrak{p}$ is an integral domain. Therefore, the kernel of the multiplication map $\mu_x : \mathcal{O}/\mathfrak{p} \to \mathcal{O}/\mathfrak{p}$ is 0 and thus μ_x is injective. Since $\mathcal{O}/\mathfrak{p}$ is finite, μ_x is also surjective so it is a bijection. Take $x^{-1} = \mu_x^{-1}(1)$. Now $xx^{-1} = x\mu_x^{-1}(1) = 1$. Hence every element of $\mathcal{O}/\mathfrak{p}$ has an inverse, and it is a field. We conclude that $\mathfrak{p}$ is maximal. □

Note that we, crucially, need that $\mathcal{O}/\mathfrak{p}$ is a finite set which is special in our setting. We eventually want to prove that any ideal is the unique product of prime ideals up to order of the factors. Towards this goal we show a simpler inclusion.

Proposition 2.14. Let $\mathcal{I}$ be a non-zero ideal of $\mathcal{O}$. Then there exists non-zero prime ideals $\mathfrak{p}_1, \ldots \mathfrak{p}_r$ of $\mathcal{O}$ such that

$$\mathfrak{p}_1\mathfrak{p}_2 \ldots \mathfrak{p}_r \subseteq \mathcal{I}$$

Proof. Let S be the set of all ideals which do not contain a such product of non-zero prime ideals. We want to prove that S is empty. Because $\mathcal{O}$ is Noetherian, S contains a maximal element, say $\mathcal{I}$. Note that $\mathcal{I}$ is maximal

with respect to the elements of S, not necessarily a maximal ideal in $\mathcal{O}$. By assumption on $\mathcal{I} \in S$, $\mathcal{I}$ is not prime. Hence there exists $\alpha, \beta \in \mathcal{I}$ such that $\alpha\beta \in \mathcal{I}$ but neither α nor β is in $\mathcal{I}$. Define two new ideals

$$\mathcal{J}_1 = \alpha\mathcal{O} + \mathcal{I} \supsetneq \mathcal{I} \quad \mathcal{J}_2 = \beta\mathcal{O} + \mathcal{I} \supsetneq \mathcal{I}.$$

Because $\alpha, \beta \notin \mathcal{I}$ we have strict inclusions. In addition we have that $\mathcal{J}_1\mathcal{J}_2 \subseteq \mathcal{I}$. Since we chose $\mathcal{I}$ to be maximal in S, neither $\mathcal{J}_1$ nor $\mathcal{J}_2$ are in S. This means that there exists $\mathfrak{p}_1, \ldots, \mathfrak{p}_r$ and $\mathfrak{q}_1, \ldots, \mathfrak{q}_s$ such that $\mathfrak{p}_1 \ldots \mathfrak{p}_r \subseteq \mathcal{J}_1$ and $\mathfrak{q}_1 \ldots \mathfrak{q}_s \subseteq \mathcal{J}_2$. But then

$$\mathfrak{p}_1 \ldots \mathfrak{p}_r\mathfrak{q}_1 \ldots \mathfrak{q}_s \subseteq \mathcal{J}_1\mathcal{J}_2 \subseteq \mathcal{I}.$$

so $\mathfrak{p}_1 \ldots \mathfrak{p}_r\mathfrak{q}_1 \ldots \mathfrak{q}_s \subseteq \mathcal{I}$, a contradiction. Therefore S does not contain a maximal element, which means that S is empty since $\mathcal{O}$ is Noetherian. □

We eventually want a group structure on ideals, but for 'normal' ideals, called integral ideals, there is not always an inverse. We therefore extend the notion of an ideal.

Definition 2.15 (Fractional Ideal)**.** Let R be an integral domain, K its field of fractions. Then an R-submodule $\mathcal{I} \subseteq K$ is a *fractional ideal* if there exists a non-zero $d \in R$ such that $d\mathcal{I} \subseteq R$.

The element $d \in R$ in the above definition can be thought of as 'cancelling' the denominators in $\mathcal{I}$. We can therefore view fractional ideals as ideals on the form $\frac{1}{d}\mathcal{J}$ for an integral ideal $\mathcal{J}$.

Letting $R = \mathbb{Z}$, $K = \mathbb{Q}$ and choosing $\mathcal{I} = \frac{1}{2}\mathbb{Z}$ we can pick the element $r = 2 \in \mathbb{Z}$ such that

$$r\mathcal{I} = 2 \cdot \left(\frac{1}{2}\mathbb{Z}\right) = \mathbb{Z} \subseteq R = \mathbb{Z}$$

and hence $\frac{1}{2}\mathbb{Z}$ is a fractional ideal in $\mathbb{Q}$.

We extend the notion of the norm to fractional ideals.

Definition 2.16 (Norm of Fractional Ideal)**.** Let $\mathcal{I}$ be a fractional ideal, i.e. there exists d' such that $d'\mathcal{I} \subseteq R$. Let d be the smallest such d'. We define the norm

$$N(\mathcal{I}) = \frac{N(d\mathcal{I})}{N(\langle d \rangle)}$$

Notice that $d\mathcal{I}$ is an ideal, and so is $\langle d \rangle$ so the norm is well defined. The norm of a fractional ideal need not be an integer, but this is still the case for integral ideals.

By the definition of an ideal we have that $\mathcal{I}\mathcal{O} = \mathcal{I}$. Now, if there exist a fractional ideal $\mathcal{J}$ such that

$$\mathcal{I}\mathcal{J} = \mathcal{O},$$

we say that $\mathcal{I}$ is invertible since $\mathcal{O}$ acts an an identity. If an ideal is invertible, the inverse has a special form. We first prove this for prime ideals.

Proposition 2.17. Let $\mathfrak{p}$ be a non-zero prime ideal of $\mathcal{O}$. Define

$$\mathfrak{p}^{-1} = \{x \in K \mid x\mathfrak{p} \subseteq \mathcal{O}\}.$$

Then we have that

1. $\mathfrak{p}^{-1}$ is a fractional ideal of $\mathcal{O}$.
2. $\mathcal{O} \subsetneq \mathfrak{p}^{-1}$
3. $\mathfrak{p}^{-1}\mathfrak{p} = \mathcal{O}$

Proof. 1. Pick a non-zero $a \in \mathfrak{p} \subseteq K$. By definition of $\mathfrak{p}^{-1}$, $a\mathfrak{p}^{-1} \subseteq \mathcal{O}$. Therefore $a\mathfrak{p}^{-1}$ is an integral ideal and $\mathfrak{p}^{-1}$ is a fractional ideal of $\mathcal{O}$

2. Clearly $\mathcal{O} \subseteq \mathfrak{p}^{-1}$. It is enough to find an element of $\mathfrak{p}^{-1}$ which is not an algebraic integer. Let $0 \neq a \in \mathfrak{p}$. By Proposition 2.14 we can choose the minimal r such that

$$\mathfrak{p}_1 \dots \mathfrak{p}_r \subseteq \langle a \rangle \quad \text{for non-zero } \mathfrak{p}_i$$

Since $\langle a \rangle \mathcal{O} \subseteq \mathfrak{p}$ and $\mathfrak{p}$ is prime, we get that $\mathfrak{p}_i \subseteq \mathfrak{p}$ for some $1 \leq i \leq r$. Without loss of generality, let $i = 1$. Now, since prime ideals are maximal by Proposition 2.13, $\mathfrak{p}_1 = \mathfrak{p}$. Removing $\mathfrak{p}_1$ from the product of prime ideals yields

$$\mathfrak{p}_2 \dots \mathfrak{p}_r \not\subseteq \langle a \rangle$$

by the minimality of the index r. We can therefore find $b \in \mathfrak{p}_2 \dots \mathfrak{p}_r$ with $b \notin \langle a \rangle$. We now claim that ba^{-1} is in $\mathfrak{p}^{-1}$ but not in $\mathcal{O}$. Since $\mathfrak{p} = \mathfrak{p}_1$, we have that $b\mathfrak{p} \subseteq \langle a \rangle \mathcal{O}$ so $ba^{-1}\mathfrak{p} \subseteq \mathcal{O}$ and $ba^{-1} \in \mathfrak{p}^{-1}$. Since $b \notin \langle a \rangle$ we have that $ba^{-1} \notin \mathcal{O}$. We have therefore found an element, namely ba^{-1}, which is in $\mathfrak{p}^{-1}$ but not in $\mathcal{O}$. We conclude that $\mathcal{O} \subsetneq \mathfrak{p}^{-1}$.

3. Here we prove that $\mathfrak{p}^{-1}$ is indeed the inverse of $\mathfrak{p}$. We have that

$$\mathfrak{p} = \mathfrak{p}\mathcal{O} \subseteq \mathfrak{p}\mathfrak{p}^{-1} = \mathfrak{p}^{-1}\mathfrak{p} \subseteq \mathcal{O}$$

Since $\mathfrak{p}$ is maximal by Proposition 2.13 we have that $\mathfrak{p}\mathfrak{p}^{-1}$ is either equal to $\mathfrak{p}$ or $\mathcal{O}$. We proceed by showing that $\mathfrak{p} = \mathfrak{p}\mathfrak{p}^{-1}$ is not possible. Assume that $\mathfrak{p} = \mathfrak{p}\mathfrak{p}^{-1}$. Let $\{\beta_1, \ldots, \beta_r\}$ be a set of generators of $\mathfrak{p}$ as an $\mathcal{O}$-module. Pick $d := ab^{-1}$, the same element as in the previous point, which is in $\mathfrak{p}^{-1}$ but not in $\mathcal{O}$. We get that

$$d\beta_i \in \mathfrak{p}^{-1}\mathfrak{p} = \mathfrak{p} \quad \text{and} \quad d\mathfrak{p} \subseteq \mathfrak{p}^{-1}\mathfrak{p} = \mathfrak{p}.$$

Now, since $d\mathfrak{p} \subseteq \mathfrak{p}$ we have

$$d\beta_i = \sum_{j=1}^{r} c_{ij}\beta_j \in \mathfrak{p}, \quad i = 1, \ldots, r$$

where $c_{ij} \in \mathcal{O}$. Equivalently

$$0 = \left(\sum_{j=1, j\neq i}^{r} c_{ij}\beta_j \right) + \beta_i(c_{ii} - d).$$

For each j we get an equation, and we can write them in matrix form as

$$\boldsymbol{C} \cdot \boldsymbol{\beta} := \begin{pmatrix} c_{11} - d & \ldots & c_{1r} \\ c_{21} & \ldots & c_{2r} \\ \vdots & \ddots & \vdots \\ c_{r1} & \ldots & c_{rr} - d \end{pmatrix} \begin{pmatrix} \beta_1 \\ \beta_2 \\ \vdots \\ \beta_r \end{pmatrix} = 0. \tag{2}$$

Therefore, the determinant of $\boldsymbol{C}$ is 0, while it is an equation of degree r in the variable d. $\mathcal{O}$ is integrally closed, and therefore $d \in \mathcal{O}$, a contradiction. We conclude that $\mathfrak{p}\mathfrak{p}^{-1} = \mathcal{O}$.

□

We proceed by proving the statement for *all* fractional ideals. This means that the set of all fractional ideals of K, denoted $\mathcal{F}_K$, forms a group under ideal multiplication where the ring $\mathcal{O}$ is the identity element. Because $\mathcal{O}$ is abelian we get that $\mathcal{I}\mathcal{J} = \mathcal{J}\mathcal{I}$, i.e. $\mathcal{F}_K$ is abelian.

Proposition 2.18. The set $\mathcal{F}_K$ of all fractional ideals of a number field K forms a group under ideal multiplication.

Proof. It is obvious that the identity element is $\mathcal{O}$. By Proposition 2.17 we have that all prime ideals are invertible. Pick therefore a non-prime integral ideal $\mathcal{I}$, with the additional property that its norm is minimal. $\mathcal{I}$ is included in a maximal ideal $\mathfrak{p}$ which, by Proposition 2.13 is also prime. Therefore

$$\mathcal{I} \subseteq \mathfrak{p}^{-1}\mathcal{I} \subseteq \mathfrak{p}^{-1}\mathfrak{p} = \mathcal{O},$$

again by Proposition 2.17. We want to show that $\mathcal{I} \neq \mathfrak{p}^{-1}\mathcal{I}$ such that the first inclusion is strict. Assume that $I = \mathfrak{p}^{-1}\mathcal{I}$. By Proposition 2.17 we can pick a $d \in \mathfrak{p}^{-1}$ but not in $\mathcal{O}$. Denote by $\{\beta_1, \ldots, \beta_r\}$ the set of generators of $\mathcal{I}$ as a $\mathcal{O}$-module. We can write

$$d\beta_i \in \mathfrak{p}^{-1}\mathcal{I} = \mathcal{I} \quad d\mathcal{I} \subseteq \mathfrak{p}^{-1}\mathcal{I} = \mathcal{I}$$

and by the same argument as in Proposition 2.17 we get that $d \in \mathcal{O}$ which contradicts our assumption. Therefore $\mathcal{I} \subsetneq \mathfrak{p}^{-1}\mathcal{I}$ and hence

$$N(\mathcal{I}) > N(\mathfrak{p}^{-1}\mathcal{I}).$$

Since we picked $\mathcal{I}$ to be the ideal of minimal norm which was *not* invertible, we get that $\mathfrak{p}^{-1}\mathcal{I}$ is invertible. Let $\mathcal{J} \in \mathcal{F}_K$ be its inverse. But this means that $\mathcal{J}\mathfrak{p}^{-1}\mathcal{I} = \mathcal{O}$, and because we, in our case, have associativity of ideal multiplication we conclude that $(J\mathfrak{p}^{-1})\mathcal{I} = \mathcal{O}$, so $\mathcal{I}$ does have an inverse.

The only thing that remains now is to show that any fracional ideal is invertible. Let $\mathcal{I}$ be a fractional ideal. We have that $\mathcal{I}$ can be written as $\frac{1}{d}\mathcal{J}$ for some integral ideal $\mathcal{J}$ and $d \in \mathcal{O}$. Therefore, since $\mathcal{J}^{-1}$ exists, $d\mathcal{J}^{-1}$ is the inverse of $\mathcal{I}$. □

Now we are ready to prove an important theorem, namely that we can factor any integral ideal in prime factors uniquely, up to permutation of ordering.

Theorem 2.19. *Any non-zero integral ideal $\mathcal{I} \subseteq \mathcal{O}$ can be written uniquely, up to ordering, as a product of prime ideals of $\mathcal{O}$.*

Proof. We start with proving existence. Let $\mathcal{I}$ be a maximal integral ideal of $\mathcal{O}$ which does not factor in prime ideals. If $\mathcal{I}$ is maximal, then it is prime by Proposition 2.13, but then it would be a product of prime ideals, namely

itself. Therefore there exists a prime (and maximal) ideal $\mathfrak{p} \supsetneq \mathcal{I}$. We then have that $\mathcal{I}\mathfrak{p}^{-1} \subsetneq \mathcal{O}$ is an integral ideal and $\mathcal{I} \subsetneq \mathcal{I}\mathfrak{p}^{-1} \subsetneq \mathcal{O}$. Now, the first inclusion is strict because if $\mathcal{I} = \mathcal{I}\mathfrak{p}^{-1}$ then $\mathfrak{p}^{-1} = \mathcal{O}$. Since we assumed $\mathcal{I}$ was a largest ideal which did not have a factorization, $\mathcal{I}\mathfrak{p}^{-1}$ must have one. Call it

$$\mathcal{I}\mathfrak{p}^{-1} = \mathfrak{p}_2 \dots \mathfrak{p}_r$$

but then

$$\mathcal{I} = \mathfrak{p}\mathfrak{p}_2 \dots \mathfrak{p}_r$$

and we get a contradiction. We conclude that any integral ideal $\mathcal{I}$ has a factorization of prime ideals.

We move on to proving the uniqueness of this factorization. Assume we have two distinct factorizations for an ideal $\mathcal{I}$

$$\mathfrak{p}_1\mathfrak{p}_2 \dots \mathfrak{p}_r = \mathcal{I} = \mathfrak{q}_1\mathfrak{q}_2 \dots \mathfrak{q}_s.$$

Let $\mathfrak{p}_1$ differ from all $\mathfrak{q}_j$. Then we pick $\alpha_j \in \mathfrak{q}_j$ but which is not in $\mathfrak{p}_1$ and consider

$$\prod \alpha_j \in \prod \mathfrak{q}_j = \mathcal{I} \subseteq \mathfrak{p}_1.$$

The last inclusion holds because $\mathfrak{p}_1$ is prime and therefore maximal. But since $\mathfrak{p}_1$ is prime and $\prod \alpha_j \in \mathfrak{p}_1$, by the definition of a prime ideal one of the α_j must lie in $\mathfrak{p}_1$, a contradiction. We conclude that $\mathfrak{p}_1$ must be equal to one of the $\mathfrak{q}_i$, say $\mathfrak{q}_1$. Then we get that

$$\mathfrak{p}_2 \dots \mathfrak{p}_r = \mathfrak{q}_2 \dots \mathfrak{q}_s$$

and, by induction, we conclude that $r = s$ and that the factorization is unique up to ordering. □

2.4 Class Group

Now let $\mathcal{P}_K$ denote the subgroup of principal fractional ideals of all fractional ideals of $\mathcal{O}$, i.e. ideal generated by one element. We construct the *class group* of K as the quotient of the group of fractional ideals on the subgroup of principal fractional ideals.

Definition 2.20 (Class Group)**.** Let $\mathcal{O}$ be the ring of integers for a field K, $\mathcal{F}_K$ the group of fractional ideals $\mathcal{P}_K \subseteq \mathcal{F}_K$ the subgroup of principal fractional ideals. The quotient group

$$\mathrm{CL}_K = \mathcal{F}_K/\mathcal{P}_K$$

is called the *class group* of K.

Observe that if $\mathcal{O}$ is a principal ideal domain then CL_K is trivial. The order of CL_K therefore measures, in some sense, in what degree the domain $\mathcal{O}$ fails to be a principal ideal domain. An important number in algebraic number theory is the *class number* of a field K. This is defined to be the size of the corresponding class group CL_K, denoted $h(K) = |\mathrm{CL}_K|$. It is desirable and conjectured that $h(K)$ is not very big. This is a long-standing open problem.

2.5 Chinese Remainder Theorem

We recall the Chinese Reminder Theorem for ideals of $\mathcal{O}$, and state some of its properties that we are going to need later.

Theorem 2.21 ([Ogg10])**.** *Let $I = \prod_{i=1}^{m} \mathcal{I}_i^{e_i}$ be the factorization of an ideal $I \subseteq \mathcal{O}$ with $\mathcal{I}_i \neq \mathcal{I}_j$ for $i \neq j$. Then there exists an isomorphism*

$$\mathcal{O}/I \to \prod_{i=1}^{m} \mathcal{O}/\mathcal{I}_i^{k_i}.$$

The two following proposition give us an way of (efficiently) compute a isomorphism between ideals $\mathcal{I}/q\mathcal{I}$ and $\mathcal{J}/q\mathcal{J}$ for any fractional ideals $\mathcal{I}$ and $\mathcal{J}$.

Proposition 2.22. Given two ideals $\mathcal{I}$ and $\mathcal{J}$ in a ring R, there exists a $t \in I$ such that $t \cdot \mathcal{I}^{-1}$ is coprime to $\mathcal{J}$.

Such a t can be computed efficiently given $\mathcal{I}$ and the prime factorization of $\mathcal{J}$[LPR10]. The following Proposition gives a way of 'canceling' the ideal $\mathcal{I}$.

Proposition 2.23. Let $\mathcal{I}$ and $\mathcal{J}$ be ideals in R and let $t \in \mathcal{I}$ be such that $t \cdot \mathcal{I}^{-1}$ is coprime to $\mathcal{J}$. Let $\mathcal{M}$ be any fractional ideal of K. The map

$$\begin{aligned} \theta_t : & K \to K \\ & u \mapsto t \cdot u \end{aligned}$$

induces an isomorphism from $\mathcal{M}/\mathcal{J}\mathcal{M}$ to $\mathcal{I}\mathcal{M}/\mathcal{I}\mathcal{J}\mathcal{M}$ as R-modules.

In particular

$$\theta_t : \mathcal{O}/\mathcal{J} \to \mathcal{I}/\mathcal{I}\mathcal{J}$$

is an isomorphism. This is achieved by choosing $\mathcal{M} = \mathcal{O}$ the multiplicative identity. Furthermore, choosing the ideal $\mathcal{J} = \langle q \rangle$ for a $q \in \mathbb{Z}$ gives a bijection

$$\theta_t : R/\langle q \rangle \to \mathcal{I}/\langle q \rangle \mathcal{I},$$

that is a bijection between the ring modulo $\langle q \rangle$ and an ideal modulo $\langle q \rangle$. This proposition can be thought of as 'canceling the ideal' $\mathcal{I}$ since the two quotients are isomorphic. Because any ring is a module over itself, considering $\mathcal{M}$ as a ring instead of a module, we get a ring *bijection* (not necessarily a ring isomorphism).

2.6 Cyclotomic Number Fields

Here we recall some facts about cyclotomic fields. A *cyclotomic number field* is $K = \mathbb{Q}(\zeta_m)$ where ζ_m is an m-th primitive root of unity, i.e. a primitive root of the polynomial $X^m - 1$. Then K is the splitting field of this polynomial. We define

$$\phi_m(X) = \prod_{(j,m)=1} (X - \zeta_n^j)$$

to be the *m-th cyclotomic polynomial*, where (j, m) denotes the greatest common divisor. We then have that

$$X^m - 1 = \prod_{d|m} \phi_d(X).$$

and that K is the splitting field of $\phi_m(X)$ over $\mathbb{Q}$. The dimension of K over $\mathbb{Q}$ is $\varphi(m)$, the Euler totient function of m. There are no real embeddings

in the canonical embedding and $\varphi(m)/2$ pairs of complex ones. In this case we get that the ring of integers is simply $\mathcal{O}_K = \mathbb{Z}[\zeta_m] \simeq \mathbb{Z}[X]/\langle X^m - 1\rangle$[Was82].

We want to use cyclotomic number fields primarily because computation is very efficient in such fields. The case where $m = 2^k$ is a power of two is currently the most widely used case, because $m/2$ is a also power of two and arithmetic modulo $\phi_m(X)$ can be done in $O(n \log n)$ time (see e.g. [LMPR08a, LPR13]). However, one needs to be careful when using cyclotomic number fields. Some lattice problems have more efficient algorithms in this case or achieve better approximated results [CDPR16, Bia15, BS16]. On the other hand, some reductions required the field to be cyclotomic [LPR10]. We come back to these points later.

3 Lattices

Lattices is a field of mathematics which was studied long before it was used in cryptography. An important reason to study lattices is that cryptographic system based on lattice theory has shown to be resistant to quantum computers, as there is yet to be found an efficient quantum algorithm for solving hard problems in lattices. In contrast, schemes based on integer factorization have been shown to be completely insecure to quantum computers[Sho97].

The most well known lattice problem is the *shortest vector problem*. In his seminal paper[Ajt96] Ajtai provided reductions between SVP and other lattice problems (see Section 4). Moreover, he showed that if there is an oracle who can solve a certain lattice problem, called *shortest integer solutions*, in the average case, then we can solve SVP *in the worst case*. This made lattices as a basis for cryptography particularly interesting, as it showed that essentially *all* instantiations of lattice problems are hard. This is not always the case for other cryptographic primitives.

Several cryptographic primitives have been based on lattices. Such schemes are very useful as we are able to achieve average-case hardness. Ajtai constructed a hash-function and proved that solving hard lattice problems reduced to finding a preimage[Ajt96], and was later proven by [GGH11] to be collision-resistant. By imposing more structure to the lattice we can make the hash-functions more effective, but this also invalidates any proof of security. As an example, [LM06] proved that if we use lattices which are also ideals in $\mathbb{Z}[X]/\langle X^n - 1\rangle$ then there exists one-way hash functions based on lattices. However, it was later shown that these schemes were not collision resistant. This is an example of how imposing more structure to increase efficiency can break the security of the scheme. In this particular case, the hash-functions were modified by [LMPR08b] to achieve collision resistance.

In this section we will introduce basic lattice theory and relate it to the algebraic number theory from Section 2. We will also define some useful lattice-quantities and introduce some tools, in particuilar some properties of the Fourier transform which relates to lattices, that we require for the reductions in Section 5.1.

3.1 Basic Lattice Theory

A lattice is a set of points in $\mathbb{R}^n$ with periodic structure. We define a lattice to be the $\mathbb{Z}$-span of a set of linearly independent vectors of $\mathbb{R}^n$.

Definition 3.1 (Lattice)**.** Let $B = \{\boldsymbol{b}_1, \boldsymbol{b}_2, \ldots, \boldsymbol{b}_n\} \in \mathbb{R}^m$ be a set of linearly independent vectors. The *lattice* Λ generated by B is

$$\Lambda = \Lambda(B) = \Big\{ \sum_{i=1}^{n} a_i \boldsymbol{b}_i \quad | \quad a_i \in \mathbb{Z} \Big\}.$$

We use the notation $\Lambda(B)$ to denote the lattice generated by the matrix B. If the basis is clear we drop is from the notation.

Therefore, a lattice is a discrete additive subgroup of $\mathbb{R}^n$ and it is not hard to see that a lattice is isomorphic to $\mathbb{Z}^n$ as $\mathbb{Z}$-modules by mapping coefficients of basis vectors to a coordinate of $\mathbb{Z}^n$. By convention, we regard the basis elements $\boldsymbol{b}_i$ as column vectors, and hence $B \in \mathbb{R}^{m \times n}$. We will interchangeably use the same notation B for a matrix of with basis elements as lattices and for a set of basis vectors. This should be clear from context. We call m the *dimension* and n the *rank* of the lattice. If we have that $m = n$ then we call Λ a *full rank lattice*. We will mostly concern ourselves with full rank lattices.

Definition 3.2 (Fundamental Parallelepiped)**.** Given a basis $B = \{\boldsymbol{b}_1, \boldsymbol{b}_2, \ldots, \boldsymbol{b}_n\}$, the *fundamental parallelepiped* is defined to be

$$P(B) = \left\{ \sum_{i=1}^{m} a_i \boldsymbol{b}_i \,|\, a_i \in [0,1)^n \right\}$$

The volume of the fundamental parallelepiped is defined to be the determinant of the lattice, i.e. $\det(\Lambda) = \sqrt{B^T B}$ and if the matrix is full rank $\det(\Lambda) = |\det(B)|$.

Lattices have a dual associated to them, and many lattice quantities depends on the structure of the dual lattice.

Definition 3.3. Let Λ be a lattice. The *dual lattice* associated to Λ, denoted $\Lambda^\vee$, is the set

$$\Lambda^\vee = \{\boldsymbol{y} \in \mathbb{R}^n \mid \langle \boldsymbol{x}, \boldsymbol{y} \rangle \in \mathbb{Z} \text{ for all } \boldsymbol{x} \in \Lambda\}.$$

It is not hard to see that $(\Lambda^\vee)^\vee = \Lambda$. It can be seen directly or by the following proposiiton.

Proposition 3.4. For a full rank lattice Λ with basis B, its dual lattice $\Lambda^\vee$ has basis $B^\vee = (B^{-1})^T$, where ^{T} denotes transposition.

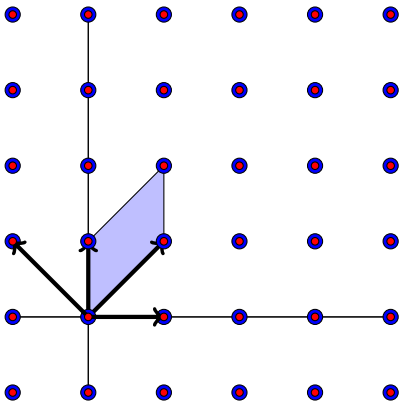

Figure 2: Two lattices. Two different bases which span the same lattice. The shaded area is the fundamental parallelepiped of one of them. Both lattices are equal to $\mathbb{Z}^2$ and are in fact dual of each other.

Proof. Let

$$B = (\boldsymbol{b}_1|\boldsymbol{b}_2|\dots|\boldsymbol{b}_n) \quad B^{-1} = (\boldsymbol{\beta}_1|\boldsymbol{\beta}_2|\dots|\boldsymbol{\beta}_n)\,.$$

By definition and linear algebra, $B^{-1}B = I$ means that

$$\langle i\text{-th row of } B^{-1}, j\text{-th column of } B\rangle = \delta_{ij}$$

and therefore

$$\langle i\text{-th column of } (B^{-1})^T, j\text{-th column of } B\rangle = \delta_{ij}.$$

By the linearity of the inner product we get the claim. □

From this proposition is follows that $\det(\Lambda) = 1/\det(\Lambda^\vee)$. We denote the basis for the dual lattice $B^\vee = \{\boldsymbol{b}_1^\vee, \boldsymbol{b}_2^\vee, \dots, \boldsymbol{b}_n^\vee\}$. This basis is unique, given a basis B, and has a special form.

Proposition 3.5. Let $\{\boldsymbol{b}_1, \dots, \boldsymbol{b}_n\}$ be a basis for the lattice Λ. Then a basis for $\Lambda^\vee$ is the unique set $\{\boldsymbol{b}_1^\vee, \dots, \boldsymbol{b}_n^\vee\}$ such that

$$\boldsymbol{b}_i \cdot \boldsymbol{b}_j^\vee = \delta_{ij}$$

where δ_{ij} is the Kronecker delta.

Proof. We first show that $\{\boldsymbol{b}_i^\vee\}$ is a basis. Let $\boldsymbol{x} = \sum_{i_1}^n c_i \boldsymbol{b}_i^\vee = 0$. Applying the inner product from the left, $\langle \boldsymbol{b}_i, - \rangle$, yields

$$\langle \boldsymbol{b}_i, \boldsymbol{x} \rangle = \sum_{i=1}^{n} \langle c_i \boldsymbol{b}_i, \boldsymbol{b}_i^\vee \rangle = c_i = \langle \boldsymbol{b}_i, 0 \rangle = 0$$

and hence $c_i = 0$ for all $i = 1, \ldots, n$ and $\{\boldsymbol{b}_i^\vee\}$ are linearly independent. From elementary linear algebra we have that any set of linearly independent vectors can be extended to a basis, but since we have n dual vectors, they are already a basis. Now, for $\boldsymbol{w} \in \mathbb{R}^n$ write it as $\boldsymbol{w} = \sum c_i \boldsymbol{b}_i^\vee$. Then $\boldsymbol{w} \cdot \boldsymbol{b}_i = c_i$, so claiming that $\boldsymbol{w}$ is in $\Lambda^\vee$ is equivalent to $c_i \in \mathbb{Z}$ for all $i = 1, \ldots n$. Therefore $\Lambda^\vee$ is the $\mathbb{Z}$-span of all the $\boldsymbol{b}_i^\vee$-s, i.e. $\Lambda^\vee$ is a lattice. □

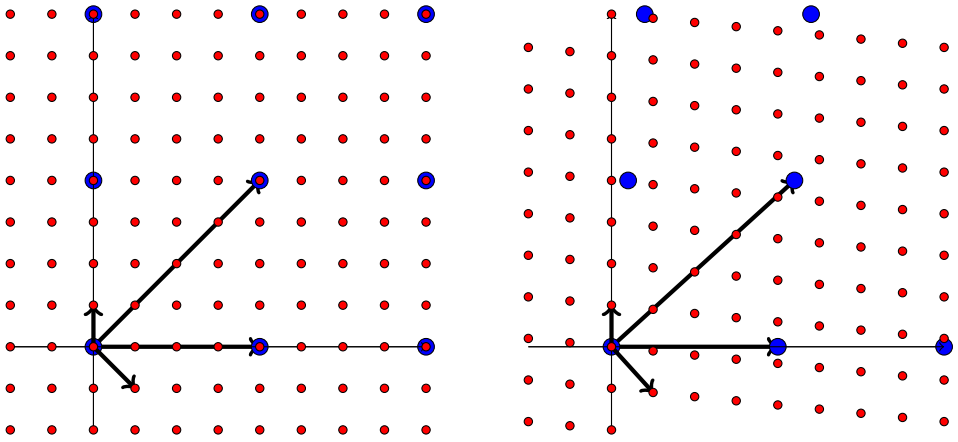

Figure 3: Two lattices (in large dots) and their dual (in small dots). The left and right large-dot lattices are spanned by $\{(2,0),(2,2)\}$ and $\{(2,0),(2.2,2)\}$ respectively. Observe that the dual is a sublattice in the left figure but not in the right. We have that $\det(\Lambda)\det(\Lambda^\vee)^{-1} = 1$ so the dual of a lattice with 'long' basis vectors will have 'short' basis vectors.

A variation of lattices is called *q-ary* lattices, which are lattices with coefficients modulo an integer.

Definition 3.6 (q-ary lattice)**.** Given a basis B of $\mathbb{R}^n$, the q-ary lattice generated by B is defined as

$$\Lambda_q(B) = \{\boldsymbol{x} = B\boldsymbol{a} \mod q \text{ for some } \boldsymbol{a} \in \mathbb{Z}^n\}$$

q-ary lattices are finite sets, as opposed to lattices in $\mathbb{R}^n$. If the lattice is full rank then a q-ary lattice contains q^n elements. They are therefore very useful in applications. In addition, many problems related to standard lattice problems have a natural modulus associated with them, making q-ary lattice an obvious candidate for reductions.

3.2 Number Field Lattice

In this section we will prove that discrete structures in K, called number field lattices, are also lattices in $\mathbb{R}^n$ under the canonical embedding σ. We do this by introducing an intermediate vector space $K \to H \to \mathbb{R}^n$.

Proposition 3.7. Consider the subspace $H \subseteq \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$ given by

$$H := \left\{(x_1, \dots x_n) \in \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2} \mid x_{s_1+s_2+j} = \overline{x_{s_1+j}}\ \forall j \in \{1. \dots, s_2\}\right\}.$$

Any discrete, additive subgroup of H is isomorphic to a lattice in $\mathbb{R}^n$.

Proof. Endowing H with the inner product $\langle \boldsymbol{x}, \boldsymbol{y} \rangle = \sum x_i \overline{y_i}$ in the ambient space $\mathbb{C}^n$ implies that H is a *real* inner product space. This means that it is isomorphic to $\mathbb{R}^n$ by an appropriate rotation. Any discrete additive subspace of H will therefore be isomorphic to a lattice in $\mathbb{R}^n$. □

Observe that H is the image of the canonical embedding: It contains real coordinates and pairs of complex coordinates which are conjugates of each other. Let us describe this isomorphism explicitly with a small example. Consider

$$H = \left\{(x, y) \in \mathbb{C}^2 \mid x = \overline{y}\right\}.$$

Now define the map $\varphi : H \to \mathbb{R}^2$ by

$$\varphi(x, y) = \left(\frac{x+y}{2}, \frac{x-y}{2i}\right) = \left(\frac{x+\overline{x}}{2}, \frac{x-\overline{x}}{2i}\right) = (\mathrm{Re}(x), \mathrm{Im}(x))$$

By the design of H, $\varphi(x, y) \in \mathbb{R}^2$. It is not hard to see that this is an isomorphism. In general we map any real coordinates of H to its own coordinate, and pair the complex in real and imaginary part. Because of the isomorphism above we consider lattices as discrete subspaces of either $\mathbb{R}^n$ or H.

We now relate lattices in H with discrete structures in $K = \mathbb{Q}(\zeta)$ because we want to make use of the nice properties of ideals from Section 2.

Definition 3.8 (Number Field Lattice)**.** Let K be a number field of degree n. A *lattice* Λ in K is the $\mathbb{Z}$-span of a $\mathbb{Q}$-basis of K

Notice the similarities with lattices in $\mathbb{R}^n$. A lattice in $\mathbb{R}^n$ is the $\mathbb{Z}$-span of a $\mathbb{R}$-basis of $\mathbb{R}^n$. Because the trace map is a bilinear form like the inner product, we define the dual lattice similarly as before.

Definition 3.9. The *dual* lattice of a number field lattice Λ is

$$\Lambda^\vee = \{x \in K \mid \operatorname{Tr}_{K/\mathbb{Q}}(xy) \in \mathbb{Z} \text{ for all } y \in \Lambda\}$$

It can be shown that the canonical embedding of a dual of a number field lattice is $\sigma(\Lambda^\vee) = \overline{\sigma(\Lambda)}^\vee$. Finally we can use our theory from Section 2.

Proposition 3.10. The ring of integers $\mathcal{O}$ is a number field lattice.

Proof. Let $\{b_1, \ldots, b_n\}$ be a $\mathbb{Z}$-basis for $\mathcal{O}$. Since $\mathbb{Q}\mathcal{O} = K$ this is also a $\mathbb{Q}$-basis for K, and it follows that $\mathcal{O}$ is a number field lattice from the definition. □

Because the canonical embedding maps $\sigma : K \to H \subseteq \mathbb{C}^n$, and a number field lattice is a discrete subspace of K we might expect that the embedding of a number field lattice under σ is a lattice in H. This is indeed the case.

Proposition 3.11. A number field lattice Λ is a lattice in H under the canonical embedding.

Proof. Let $B = \{u_1, \ldots, u_n\}$ be a $\mathbb{Q}$-basis for K. Consider

$$\sigma(B) := \{\sigma(u_1), \ldots, \sigma(u_n)\},$$

the image of the basis vectors (columns of B) under the canonical embedding. To see that the set $\sigma(B)$ is linearly independent assume that

$$\sigma(u_1) = \sum_{i=2}^{n} \alpha_i \sigma(u_i) \quad \alpha_i \in \mathbb{R}.$$

Since σ is linear and keeps $\alpha_i \in \mathbb{R}$ fixed we have

$$\sigma(u_1) = \sigma\left(\sum_{i=2}^{n} \alpha_i u_i\right).$$

but since σ is injective, this means that

$$u_1 = \sum_{i=2}^{n} \alpha_i u_i$$

which is impossible because $\{u_i\}$ is a basis. We conclude that the set $\{\sigma(u_i)\}$ is linearly independent.

Now by the definition of H we have that $\sigma(u_i) \in H$. Since $H \simeq \mathbb{R}^n$ and we have n linearly independent vectors $\sigma(u_i)$ we conclude that the $\mathbb{Z}$-span of $\sigma(B)$ is a lattice in H. □

This means that any number field lattice is a lattice in H under the canonical embedding, which in turn is isomorphic to a lattice in $\mathbb{R}^n$. In particular, the ring of integers $\mathcal{O}$ and its ideals are lattices in $\mathbb{R}^n$ in this way. We can now permit ourselves to talk about various lattice quantities in all these lattices. For instance, the shortest vector of $\mathcal{O}$ is defined to be the shortest vector in the corresponding lattice in $\mathbb{R}^n$.

Let us see how the ring of integers of a cyclotomic field is a lattice in $\mathbb{R}^n$. Let ζ be the 3-rd root of unity, $\zeta^3 = 1$. Then the ring of integers in $K = \mathbb{Q}[\zeta]$ is equal to $\mathcal{O} = \mathbb{Z}[\zeta] = \text{span}_{\mathbb{Z}}\{1, \zeta\}$. We have the two embeddings $\sigma_i : K \to \mathbb{C}$

$$\begin{aligned} \sigma_1 &: \zeta \mapsto \zeta \\ \sigma_2 &: \zeta \mapsto \zeta^2 \end{aligned}$$

with $\sigma_1 = \overline{\sigma_2}$. The image of the basis for $\mathcal{O}$ becomes

$$\sigma(\mathcal{O}) = \text{span}_{\mathbb{Z}}\{\sigma(1), \sigma(\zeta)\} = \text{span}_{\mathbb{Z}}\{\begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} \zeta \\ \zeta^2 \end{pmatrix}\} \in H.$$

By now applying the isomorphism described in Proposition 3.7 we finally get

$$\begin{aligned} \phi(1,1) &= (1,0) \\ \phi(\zeta, \zeta^2) &= (-1/2, 1/2) \end{aligned}$$

which is a basis for $\mathbb{R}^2$. The $\mathbb{Z}$-span of this basis is therefore a lattice, shown in Figure 4.

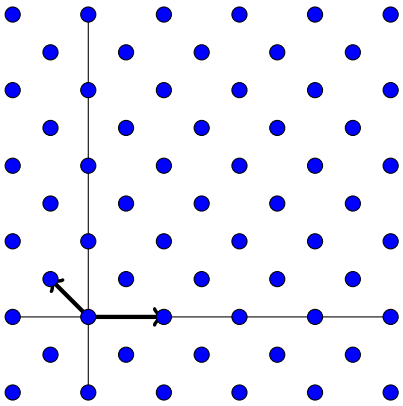

Figure 4: Embedding the number field lattice $\mathcal{O} = \mathbb{Z}[\zeta_3]$ into $\mathbb{R}^2$.

3.3 Ideal Lattices

A problem with cryptosystems based on lattices is that the key size is often large, making lattice based cryptography inefficient. However, by using more structured lattice we can hope to reduce the key-size. See an example of this in [RS10]. By imposing more structure on the lattice we invalidate the proof of security, and new analysis is required.

A widely used such lattice is an *ideal lattice*. This is a lattice which comes from an ideal of the ring of integers.

Definition 3.12 (Ideal Lattice)**.** Let $\sigma : K \to \mathbb{C}^n$ be the canonical embedding. A lattice Λ in H is an *ideal lattice* if there exists an ideal $\mathcal{I} \subseteq \mathcal{O}$ such that $\sigma(\mathcal{I}) \simeq \Lambda$.

By this definition an ideal lattice corresponds to an ideal in $\mathcal{O}$. We also have the converse: any ideal in $\mathcal{O}$ corresponds to a *full rank* lattice in H.

Proposition 3.13. Every ideal $\mathcal{I}$ of $\mathcal{O}$ is corresponds to a full rank lattice in H.

Proof. Let $\mathcal{I} = \text{span}_{\mathbb{Z}}\{b_1, \ldots, b_n\}$. Embedding $\mathcal{I}$ with σ yields

$$\sigma(\mathcal{I}) = \{\sigma(b_1), \ldots, \sigma(b_n)\}$$

which is a linearly independent set because σ is an injection and $\{b_i\}$ is a basis. Since there are n vectors we conclude that it forms a basis for H and hence $\mathbb{Z} \cdot \sigma(\mathcal{I})$ is a lattice in H of rank n. □

We will make use of a bound of short vectors in ideal lattices, namely that for any ideal $\mathcal{I} \subseteq \mathcal{O}$, the length $\lambda_1(\mathcal{I})$ of the shortest vector is bounded.

Proposition 3.14. For an ideal lattice Λ, $\lambda_1(\mathcal{I}) \geq \sqrt{n} \cdot \mathrm{N}(\mathcal{I})^{1/n}$

and in particular for the ring of integers itself

$$\lambda_1(\mathcal{O}) \geq \sqrt{n}.$$

See Section 4 for the definition and discussion of $\lambda_1(\mathcal{O})$. Because the length of a lattice point in a number field lattice is defined as the length of the embedded vector in $\mathbb{R}^n$, any ideal lattice satisfies this bound. Note that for any general lattice in $\mathbb{R}^n$, we are free to scale the basis elements to create a new lattice of the same rank with shorter vectors. Therefore the above bound does not hold in general, but is nevertheless useful when we are in the ideal lattice case.

3.4 Fourier Transform

We use the Fourier transform as a tool in the reductions in Section 5. For a discrete distribution D, the Fourier transform of $\boldsymbol{x} \in \mathbb{R}^n$ is defined to be

$$f(\boldsymbol{x}) := \sum_{\boldsymbol{y} \ D} D(\boldsymbol{y}) \exp(2\pi i \langle \boldsymbol{x}, \boldsymbol{y} \rangle) = \mathrm{E}[D(\boldsymbol{y}) \exp(2\pi i \langle \boldsymbol{x}, \boldsymbol{y} \rangle)]$$

where the second equality is from the definition of the expectation. Given enough samples $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_N$it can be shown that we can approximate the Fourier transform by

$$f(\boldsymbol{x}) \approx \frac{1}{N} \sum_{j=1}^{N} \exp(2\pi i \langle \boldsymbol{x}, \boldsymbol{y}_j \rangle)$$

By using the same idea for samples from $D_{\Lambda,r}$ we can approximate the Fourier transform for $D_{\Lambda,r}$ by sampling from $D_{\Lambda,r}$. Denote the Fourier transform of $D_{\Lambda,r}$ by $f_{1/r}$. If $1/r$ is sufficiently small it can be shown that we have the approximation

$$f_{1/r}(\boldsymbol{x}) \approx \exp\left(-\pi(r \cdot \mathrm{dist}(\Lambda^\vee, \boldsymbol{x}))^2\right).$$

Observe that the Fourier transform depends on the distance to the *dual* lattice $\Lambda^\vee$. For $\boldsymbol{x} \in \Lambda^\vee$ we get that $f_{1/r}(\boldsymbol{x}) \approx 1$. Indeed, for any $\boldsymbol{x} \in \Lambda^\vee$ we get that

$$f(\boldsymbol{x}) = \sum_{\boldsymbol{y} \ D} D(\boldsymbol{y}) \exp(2\pi i \langle \boldsymbol{x}, \boldsymbol{y} \rangle) = \sum_{\boldsymbol{y} \ D} D(\boldsymbol{y}) = 1$$

because $\langle \boldsymbol{x}, \boldsymbol{y} \rangle \in \mathbb{Z}$ when $\boldsymbol{x} \in \Lambda^\vee$.

4 Interesting Lattice Problems

The most well known lattice problem is the *shortest vector problem* (SVP), where we are given a basis for a lattice and are required to output the shortest vector in the lattice. Closely related is the *closest vector problem* (CVP) where we are required to find the closest lattice point to an arbitrary point in $\mathbb{R}^n$. In this section we will present standard lattice problems, together with simple attacks or references to such attacks. Then we will show reductions between the problems to build an understanding of how the problems are related.

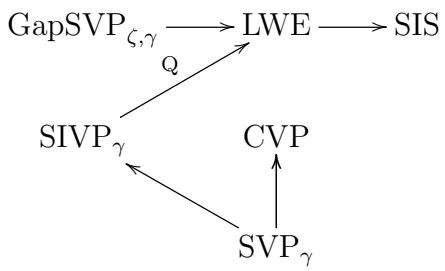

Figure 5: Problem Tree. The arrows indicate directions of reduction. Quantum reductions are indicated with a Q. This is not the complete picture, but illustrates what we do in this thesis.

In this section we will present many of the standard lattice problems. The description of the problems are threefold. Firstly we describe define the problem, secondly we describe standard attacks to illustrate how well we can solve them and thirdly we provide reductions between the problems to show how they relate to each other.

4.1 Shortest Vector Problem

The most well known computational problem in lattice theory is the shortest vector problem. We start by defining a very useful lattice quantity which measures the shortest length of a set of linearly independent vectors in the lattice.

Definition 4.1 (Successive Minima)**.** Let Λ be a lattice and $\mathcal{B}(0, r)$ the ball of radius r centered at 0. The *n-th succesive minimum*, denoted, λ_n, is the smallest r such that $\mathcal{B}(0, r) \cap \Lambda$ contains n linearly independent elements.

Note that $\lambda_1(\Lambda)$ is the length of the shortest vector in the lattice. On the other hand, $\lambda_2(\Lambda)$ is not necessarily the length of the second shortest vector because such a vector might be linearly dependent on the shortest.

Definition 4.2 (SVP)**.** Let B be a basis for the lattice Λ. The *shortest vector problem* is, given B, to output a vector of length $\lambda_1(\Lambda)$.

Definition 4.3 (SIVP)**.** Let B be a basis for the lattice Λ. The *shortest independent vector problem* (SIVP) is to output n linearly independent vectors $\{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n\} \in \Lambda$ such that $||\boldsymbol{x}_i|| \leq \lambda_n(\Lambda)$ for $i = 1, \ldots, n$.

If the lattice Λ is clear from context we denote $\lambda_n(\Lambda)$ by λ_n. It is easy to see that $0 < \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$. How hard these problems are depends to a large extent on which basis describes the lattice we are given as input. For instance, consider a lattice spanned by an orthogonal basis. It is not hard to see that the shortest vector in such a lattice would be one of the basis vectors. Finding the successive minima is equivalent to sorting the basis by length, which can be done in $O(n \log n)$ time. In the general case, given a basis B for a lattice we can try to make B orthogonal, or as orthogonal as possible, and solve SVP in this way. However, there is no known polynomial algorithm that solves neither SIVP nor SVP in this way. If we instead are required to output approximations of short vectors up to a factor γ then it can be done.

Definition 4.4 (SVP_γ)**.** Given a basis B for a lattice Λ, output a non-zero vector of length at most $\gamma \cdot \lambda_1(\Lambda)$ for $\gamma = \gamma(n)$.

Definition 4.5 (SIVP_γ)**.** Given a basis B of a lattice λ, output n linearly independent vectors *all* of length at most $\gamma\lambda_n$ for $\gamma = \gamma(n)$.

By setting $\gamma = 1$ we get the exact versions of SIVP. For these relaxed versions we have some interesting attacks.

4.1.1 Attack

There is a polynomial time algorithms which achieves exponential γ: Use the revered δ-LLL algorithm to recover a $O(2^n)$-approximate set if short, linearly

independent elements in $\tilde{O}(n^6)$ time where n is the lattice dimension[KLWLL82]. This lattice reduction algorithm calculates the so-called LLL-reduced basis, a 'nearly orthogonal' basis of elements sorted length. We illustrate the idea in the case $n = 2$.

Assume we have a basis $B = \{\boldsymbol{b}_1, \boldsymbol{b}_2\}$ for a lattice in $\mathbb{R}^2$ and that $||\boldsymbol{b}_1|| \leq ||\boldsymbol{b}_2||$. If this is not the case we just flip their order. Now compute the Gram-Schmidt coefficient $\mu_{21} = \langle \boldsymbol{b}_2, \boldsymbol{b}_1 \rangle / ||\boldsymbol{b}_1||^2$ and reduce the length of the longest vector by computing $\boldsymbol{b}_2 = \boldsymbol{b}_2 - \lfloor \mu_{21} \rceil \boldsymbol{b}_1$. Repeat this step until we are no longer able to reduce the basis. The same idea is used in higher dimensions but some more care needs to be taken[KLWLL82].

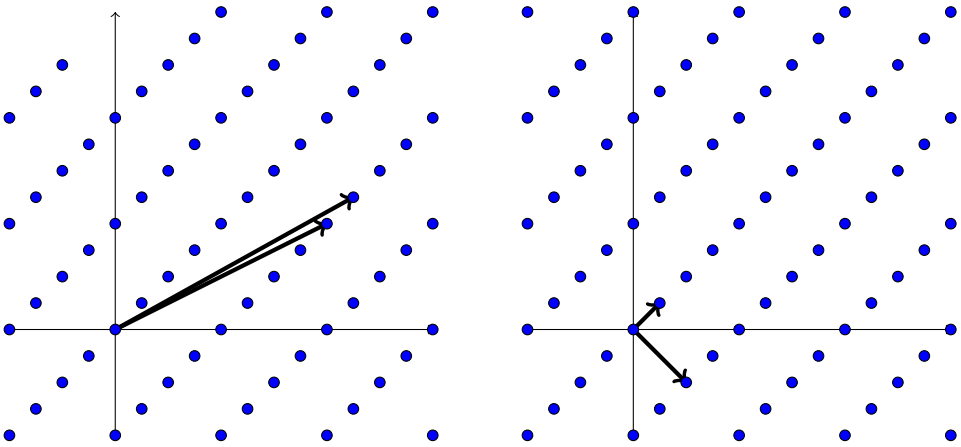

Figure 6: Two equal lattices Λ with different bases. To the right is the LLL-reduced basis.

In Figure 6 we see a lattice generated by a basis and its LLL-reduced counterpart. Notice that, in this case, we have an exact ($\gamma = 1$) solution of the $\text{SVP}_\gamma$and SIVP_γ, namely the LLL-reduced basis vectors. This is because there exists an orthogonal basis for this lattice. In general however, we expect an exponential approximation.

The hardness of SVP_γ(and SIVP_γ) for polynomial γ is what we base many of the cryptographic schemes on. Indeed, Ajtai showed that the exact SVP is NP-hard under randomized reductions, increasing our confidence that SVP is a hard problem[Ajt98]. There exists no polynomial time algorithm, neither classical nor quantum, for SIVP that achieves a polynomial approximation factor γ. We therefore conjecture that SVP and SIVP are hard, even for quantum algorithms.

4.1.2 Variants of SVP

We briefly mention the decision version of SVP which is to determine whether λ_1 is smaller than a given d or larger than $\gamma \cdot d$. If neither of these are true any answer is accepted.

Definition 4.6 (GapSVP_γ)**.** Given a basis for a lattice Λ and a real $d > 0$, determine whether $\lambda_1(\Lambda) \leq d$ or $\lambda_1(\Lambda) > \gamma \cdot d$. It is a YES instance in the former case and NO instance in the latter.

We have a trivial reduction from GapSVP_γ to SVP: If we are able to recover the shortest vector in a lattice then it is easy to check its length. There is no known reduction in the other direction.

4.2 Closest Vector Problem

Closely related to SVP is the *closest vector problem*. We denote by $\text{dist}(\Lambda, \boldsymbol{y})$ the distance from an arbitrary $\boldsymbol{y} \in \mathbb{R}^n$ to the lattice Λ.

Definition 4.7 (CVP_γ)**.** Given a basis B of a lattice Λ and a vector $\boldsymbol{y} \in \mathbb{R}^n$, output a non-zero vector $\boldsymbol{x}$ such that $||\boldsymbol{y} - \boldsymbol{x}|| \leq \gamma \cdot \text{dist}(\Lambda, \boldsymbol{y})$.

Again by choosing $\gamma = 1$ we get an exact version of CVP. There is a polynomial time algorithm for CVP_γ using a LLL-reduced basis and an algorithm often referred to as "Babai's nearest plane algorithm' due to Babai[Bab86]. This algorithm achieves an exponential approximation factor. We omit any discussion of this algorithm.

4.2.1 Hardness

We expect that there is a reduction from CVP to SVP: Use a CVP oracle to output the closest lattice point to 0. Since $0 \in \Lambda$ this algorithm would just output 0 which is not a solution to SVP. We can construct an algorithm with some slight modifications. Let B be a basis of the lattice Λ. Let $\Lambda(B^{(i)})$ denote the lattice generated by $B^{(i)} = \{\boldsymbol{b_1}, \ldots, 2\boldsymbol{b}_i, \ldots, \boldsymbol{b}_n\}$, i.e. the same lattice but with its i-th basis vector scaled by a factor 2. It is easy to see that both

$$\boldsymbol{u} \in \Lambda(B^{(i)}) \quad \Rightarrow \quad \boldsymbol{v} = \boldsymbol{u} - \boldsymbol{b}_i \in \Lambda(B) \tag{3}$$

and

$$\boldsymbol{v} = \sum_{i=1}^{n} \alpha_i \cdot \boldsymbol{b}_i \text{ such that } \alpha_j \text{ is odd for } j \quad \Rightarrow \quad \boldsymbol{u} = \boldsymbol{v} - \boldsymbol{b}_j \in \Lambda(B^{(j)}). \quad (4)$$

We show this reduction in the case where $\gamma = 1$ for simplicity.

Proposition 4.8. There is a polynomial time reduction from SVP to CVP.

```
1 SVP^CVP(B)
2 for i = 1 to n do
3 |   B^(i) ← {b_1, ..., 2b_i, ..., b_n}
4 |   x_i ← CVP(B^(i), b_i)
5 end
6 return min{x_i − b_i}
```

Figure 7: SVP to CVP reduction.

Proof. The algorithm is shown in Figure 7. Suppose that $\boldsymbol{v} = \sum \alpha_i \cdot \boldsymbol{b}_i$ is the shortest vector in $\Lambda(B)$. If every coordinate of $\boldsymbol{v}$ is even then $\boldsymbol{v}' = \boldsymbol{v}/2 \in \Lambda(B)$ and is shorter. Hence at least one coordinate is odd, say α_1. In the lattice $\Lambda(B^{(1)})$ we have that $\boldsymbol{v} + \boldsymbol{b}_1$ is the closest vector to $\boldsymbol{b}_1$ by (4). $\mathrm{CVP}(B^{(1)}, \boldsymbol{b}_1)$ will therefore output $\boldsymbol{w} = \boldsymbol{v} + \boldsymbol{b}_1$, which, by (3) is a vector in $\Lambda(B)$. Since we know the basis elements we can easily recover $\boldsymbol{v} = \boldsymbol{w} - \boldsymbol{b}_1$. □

This means that SVP is at least as hard as CVP.

4.2.2 Variant of CVP

There is a variant of the CVP_γ is called *bounded distance decoding*. In this problem, the challenge $\boldsymbol{y}$ is sufficiently close to the lattice such that the solution is unique.

Definition 4.9 (Bounded Distance Decoding Problem (BDD_γ))**.** Let B be a basis for a lattice Λ. Given a point $\boldsymbol{y} \in \mathbb{R}^n$ such that $(\gamma + 1) \cdot \mathrm{dist}(\Lambda, \boldsymbol{y}) < \lambda_1(\Lambda)$, output the unique $\boldsymbol{x} \in \Lambda$ closest to $\boldsymbol{y}$.

In other words, given a $\boldsymbol{x} = \boldsymbol{y} + \boldsymbol{e}$ for a $\boldsymbol{y} \in \Lambda$ and a bounded 'perturbation' $\boldsymbol{e}$, recover $\boldsymbol{x}$. The solution is unique because if $||\boldsymbol{y} - \boldsymbol{x}|| < \lambda_1(\Lambda)/(\gamma+1)$ then for any $\boldsymbol{z} \in \Lambda \backslash \{\boldsymbol{x}\}$

$$||\boldsymbol{y} - \boldsymbol{z}|| \geq ||\boldsymbol{x} - \boldsymbol{z}|| - ||\boldsymbol{y} - \boldsymbol{x}|| > \gamma \cdot \lambda_1(\Lambda)/(\gamma + 1) > \gamma \cdot \mathrm{dist}(\Lambda, \boldsymbol{y})$$

Clearly we have a reduction from BDD_γ to CVP_γ . Solving CVP_γ can only be easier if we know that the solution is close to the challenge point.

A very simple algorithm to solve BDD is, on input basis B and vector $\boldsymbol{t} \in \mathbb{R}$, is called the round off algorithm and it simply outputs $B \cdot \lfloor (B^\vee)^T \cdot \boldsymbol{y} \rceil$.

Proposition 4.10. Let Λ be a lattice in $\mathbb{R}^n$ with basis B, and let $\boldsymbol{y} = \boldsymbol{x} + \boldsymbol{e} \in \mathbb{R}^n$ be such that $\langle \boldsymbol{b}_j^\vee, \boldsymbol{e} \rangle \in \left[-\frac{1}{2}, \frac{1}{2}\right)$ for all j for $\boldsymbol{v} \in \Lambda$. Then the round-off

$$B \cdot \lfloor (B^\vee)^T \cdot \boldsymbol{y} \rceil = \boldsymbol{x}$$

returns the desired element $\boldsymbol{x}$.

Proof. Recall that $B^{-1} = (B^\vee)^T$. Because $\boldsymbol{x}$ is in the lattice have have that $\boldsymbol{x} = B\boldsymbol{z}$ for some integer vector $\boldsymbol{z} \in \mathbb{Z}^n$. Compute

$$\lfloor (B^\vee)^T \cdot \boldsymbol{y} \rceil = \lfloor \boldsymbol{z} + (B^\vee)^T \cdot \boldsymbol{e} \rceil.$$

since $B^{-1} = (B^\vee)^T$. On the assumption on $\langle \boldsymbol{b}_j, \boldsymbol{e} \rangle$ is small, we get that

$$\lfloor (B^\vee)^T \boldsymbol{y} \rceil = \lfloor \boldsymbol{z} + (B^\vee)^T \cdot \boldsymbol{e} \rceil = \boldsymbol{z}$$

because $\boldsymbol{z}$ is an integer vector and each coordinate of $(B^\vee)^T \boldsymbol{e}$ is in $[-1/2, 1/2)$. Since $\boldsymbol{z}$ was the vector of integer coefficients of $\boldsymbol{x}$ we compute $B\boldsymbol{z} = \boldsymbol{x}$. □

The assumption on $\langle \boldsymbol{b}_j^\vee, \boldsymbol{e} \rangle$ is quite strict so this attack has limited uses. It is nevertheless useful in some contexts, e.g. in [CDPR16]. Before we move on to more, less standard lattice problems we need a short intermesso into the world of probability.

4.3 Probability

In this subsection we define an important distribution for the hardness theorems, reductions and applications. We also state some additional properties from statistics, such as the statistical distance between distributions and how uniform distributions behave under bijective maps.

4.3.1 Statistical Tools

We often need to guarantee that samples drawn in a particular way come from a desired distribution. However, by allowing samples to come from a distribution which is *essentially* the same we give ourselves a bit more freedom. We formalize this by introducing a metric on the set of probability spaces.

Definition 4.11. For two distributions P and Q on the sample set X, the *statistical distance* between P and Q is

$$\Delta(P, Q) := \sup_{A \subseteq X} |P(A) - Q(A)|$$

We say the two distributions are *statistically indistinguishable* if $\Delta(P, Q) < \varepsilon$ for some negligible ε.

In many of the attacks we describe, we require that uniform distributions remain uniform under bijective maps.

Proposition 4.12. Let $f : A \to B$ be a map between sets A and B. Assume $|f^{-1}(b)| = n \in \mathbb{N}$ is constant for all $b \in f(A)$. If $a \xleftarrow{r} A$ then the distribution of $f(a)$ is uniform.

Proof. Assume, for simplicity, that A and B are finite, such that each element $a \in A$ is assigned probability $1/|A|$. Since every preimage of elements B contains n elements in A, $f(a)$ is assigned probability $n/|A|$. We conclude that $f(a)$ is uniform on $f(A) \subseteq B$. □

In particular, if f is a bijection then $|f^{-1}(b)| = 1$ and $f(A) = B$ so the distribution of $f(a)$ is uniform (with the same probability assigned to each element) on all of B.

4.3.2 Gaussian Distributions

A particular distribution we are going to discuss is the Gaussian distribution of width r, D_r. We first define the Gaussian function

$$\rho_r(\boldsymbol{x}) := \exp(-\pi ||\boldsymbol{x}||^2 / r^2).$$

It it not hard to see that the measure of ρ_r over $\mathbb{R}^n$ is r^n. The Gaussian distribution D_r should be the distribution proportional to ρ_r, which is obtained by scaling: $D_r(\boldsymbol{x}) = \rho_r(\boldsymbol{x})/r^n$. We similarly define the Gaussian $D_{\Lambda,r}$ where samples are lattice points $\boldsymbol{x} \in \Lambda$.

$$D_{\Lambda,r}(\boldsymbol{x}) = \frac{\rho_r(\boldsymbol{x})}{\rho_r(\Lambda)} \quad \boldsymbol{x} \in \Lambda.$$

Even though we can sample from the Gaussian D_r easily, sampling its discrete counterpart $D_{\Lambda,r}$ is not always trivial.

Definition 4.13 ($\mathrm{DGS}_{\Lambda,r}$)**.** The *discrete gaussian sampling* problem (DGS) with width r over the lattice Λ is to output a sample from $D_{\Lambda,r}$.

However, for large enough r we can sample from $D_{\Lambda,r}$ is classical polynomial time. We omit the proof.

Proposition 4.14. There exists an efficient algorithm that, given any n-dimensional lattice Λ and $r \geq 2^{2n}\lambda_n(\Lambda)$, outputs a sample from a distribution that is within distance $2^{-\Omega(n)}$ of $D_{\Lambda,r}$.

Proof. See [Reg09, Lemma 3.2]. □

The Gaussian distribution is used because it gives provable security ([Pei16, Reg09, LPR10]) when the errors are sampled from a correctly chosen $D_{\Lambda,r}$. Because of the exponential decay of $D_{\Lambda,r}$, a sample from a Gaussian of narrow width will, with high probability, be short. We show how this relates to standard lattice problems in 4.5.

Proposition 4.15. For any n-dimensional lattice Λ, any sample from$D_{\Lambda,r}$ has Euclidean norm at most $r\sqrt{n}$ except with probability 2^{-2n}.

Proof. See [Cai03, Lemma 1.5]. This is not surprising because of the exponential decay of the distribution. □

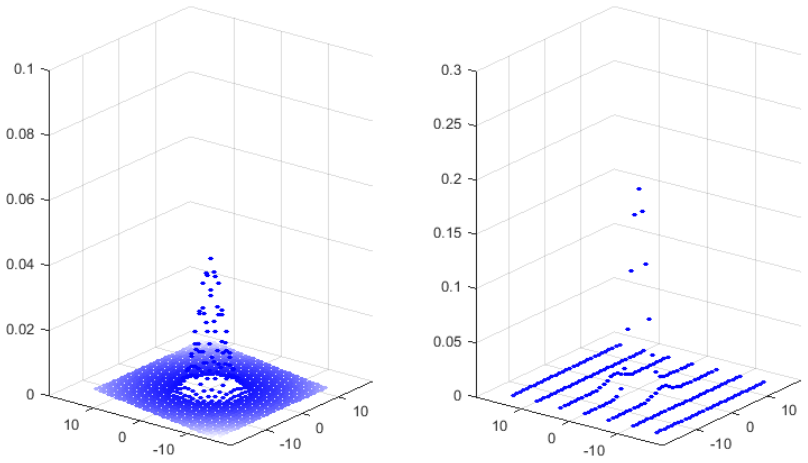

Figure 8: Gaussians $D_{\Lambda,r}$ with the same width $r = 5$ over a lattice and its dual. Here z-value indicates probability.

4.4 Lattice Quantities

Micciancio and Regev[RM07] defined a lattice parameter, called the smoothing parameter, which determines the smallest r such that for larger r, the statistical distance between $D_{\Lambda,r}$ and the uniform distribution on the lattice is negligible.

Definition 4.16. For a lattice Λ and an $\varepsilon > 0$ the *smoothing parameter* $\eta_\varepsilon(\Lambda)$ is the smallest $r > 0$ such that

$$\rho_{1/r}(\Lambda^\vee \backslash \{0\}) \leq \varepsilon.$$

The reason for why the smoothing parameter for a lattice depends on the image of (most of) its *dual* under ρ is because we have the relation[RM07], called the Poissons summation formula,

$$\rho(\Lambda) = \det(\Lambda^\vee)\hat{\rho}(\Lambda^\vee),$$

where $\hat{\rho}$ denotes the Fourier transform. This relates the two quantities. The smoothing parameter is important in analyzing security, as it gives us a measure on whether an attacker sees D_r or a uniform one. This is illustrated in the following proposition.

Proposition 4.17. For any lattice Λ, $\varepsilon > 0$ and a width $r > \eta_\varepsilon(\Lambda)$, the statistical distance between the $D_r \mod \Lambda$ and the uniform distribution $U(\Lambda)$ on Λ is less than $\varepsilon/2$.

$$\Delta(D_r \mod \Lambda, U(\Lambda) \mod \Lambda) \leq \varepsilon/2$$

Proof. See [RM07, Lemma 4.1] □

For samples from D_r with width r which exceeds the smoothing parameter, an attacker cannot distinguish these samples from uniform ones. Additionally, for an encryption scheme which on input $\boldsymbol{x} \in \Lambda$ encrypts it by adding some noise $\boldsymbol{x} + \boldsymbol{e}$ with $\boldsymbol{e} \leftarrow D_r$, then decryption becomes impossible if r exceeds $\eta_\varepsilon(\Lambda)$. Choosing the correct r is therefore essential. We state some bounds for the smoothing parameter, connecting it to previously defined lattice quantities.

Proposition 4.18. For an n-dimensional lattice Λ we have that

$$\eta_\varepsilon(\Lambda) \leq \sqrt{\log(n/\varepsilon)}\lambda_n(\Lambda) \qquad 0 < \varepsilon < 1$$

In particular, using the notation $\omega(f)$ to represent a function growing more rapidly than f, we have that $\eta_\varepsilon(\Lambda) \leq \omega(\sqrt{\log n})\lambda_n(\Lambda)$. From [RM07] we get a similar upper bound of the smoothing parameter but depending inversely on the shortest vector of the dual ideal.

Proposition 4.19. For a lattice Λ of dimension n we have that

$$\eta_\varepsilon(\Lambda) \leq \sqrt{n}/\lambda_1(\Lambda^\vee)$$

Consider an encryption scheme where we encrypt by adding Gaussian noise to a lattice point. If we add too much noise, then the distribution of the resulting ciphertext will be essentially uniform and we are no longer able to decrypt correctly. The bound above therefore tells us how much noise we can add to a message and still be able to decrypt correctly, and it depends inversely on the shortest vector in the dual lattice. Minimizing $\lambda_1(\Lambda^\vee)$ therefore permits us to add a large amount of noise while still being able to decrypt.

4.5 More Lattice Problems

Recall that from Proposition 4.15, a sample from $D_{\Lambda,r}$ will have length $\geq r\sqrt{n}$ with high probability. By choosing r sufficiently small and sampling from $D_{\Lambda,r}$ we expect to get a small lattice vector. Proposition 4.18 gives us that

$$\eta_\varepsilon(\Lambda) \leq \omega(\sqrt{\log n})\lambda_n(\Lambda).$$

Now choosing

$$\gamma = \eta_\varepsilon(\Lambda) \cdot \tilde{O}(1/\alpha) \leq \omega(\sqrt{\log n})\lambda_n(\Lambda) \cdot \tilde{O}(1/\alpha)$$

and observing that

$$\gamma\sqrt{n} = \tilde{O}(\sqrt{n}/\alpha)\lambda_n(\Lambda)$$

means that a sample from $D_{\Lambda,\gamma}$ will have length $\gamma\sqrt{n} = \tilde{O}(\sqrt{n}/\alpha)\lambda_n(\Lambda)$ with high probability. Such a sample gives us a $\tilde{O}(\sqrt{n}/\alpha)$-approximated shortest vector because $\lambda_1 \leq \lambda_n$. This gives us a reduction from $\text{SVP}_{\sqrt{n}/\alpha}$ to DGS_γ. By drawing multiple samples $\boldsymbol{v}_i \leftarrow D_{\Lambda,\gamma}$ we expect to get linearly independent, short samples when we draw more than n samples. This is similar to how choosing random integer vectors $\boldsymbol{a}_i \xleftarrow{r} \mathbb{Z}^n$ quickly result in a linearly independent set $\{\boldsymbol{a}_i\}$ when we draw slightly more than n samples. A (quit technical) proof of this is found in [Reg09]. Because we have conjectured that SVP is hard then so is $\text{DGS}_{\Lambda,r}$ *for narrow with* r.

For completeness we define the *shortest integer solution* (SIS) problem.

Definition 4.20 (Shortest Integer Solution (SIS))**.** Given $\boldsymbol{a}_i \in \mathbb{Z}_q^n$ and $\beta > 0$, find integers $\boldsymbol{z} \in \mathbb{Z}^n$ such that

$$z_1\boldsymbol{a}_1 + z_2\boldsymbol{a}_2 + \cdots + z_n\boldsymbol{a}_n = 0 \quad \in \mathbb{Z}_q^n \tag{5}$$

and $||\boldsymbol{z}|| < \beta$. In other words, $A\boldsymbol{z} = 0 \bmod q$ for a small integer vector $\boldsymbol{z}$.

This is equivalent to finding an appropriate integer vector in the lattice

$$\Lambda_q(A)^\perp := \{\boldsymbol{x} \in \mathbb{Z} \mid A\boldsymbol{x} = 0 \mod q\}$$

Note that even though we need to find small $\boldsymbol{z}$, this does not immediately mean that the lattice vector $A\boldsymbol{z}$ is small. A small lattice vector can have

large coefficients. To illustrate a reduction from decision-LWE to SIS, a SIS oracle can recover a $\boldsymbol{w} \in \frac{1}{q}\Lambda(A)^{\perp}$ for A the matrix from LWE and modulus q. We can then check what the distribution of $\langle \boldsymbol{b}, \boldsymbol{w} \rangle \mod \mathbb{Z}$ is for a decision-LWE samples $\boldsymbol{b}$. See [RS10] for details. As mentioned in the start of this section, Ajtai[Ajt96] showed a reduction from worst-case approximated lattice problems to average case SIS.

5 Learning With Errors

In Section 4 we discussed the relationship between many standard lattice problems. In this section, we will see how these problems relate to the *learning with errors* (LWE) problem. In particular, we have a quantum reduction from SIVP_γ to LWE due to Regev[Reg09]. It uses an iterative step to produce samples from more and more narrow Gaussians $D_{\Lambda,r}$ which means that, eventually, one of the samples is small by Proposition 4.15. The first part consists on solving a BDD instance on the dual lattice $\Lambda^\vee$ by making use of an LWE oracle. This is done by scaling samples from $D_{\Lambda,r}$ by $1/q$ and approximating the Fourier transform of the scaled distribution $\frac{1}{q}D_{\Lambda,r}$. The Fourier transform of a vector $\boldsymbol{x}$ from $\frac{1}{q}D_{\Lambda,r}$ depends, in some way, on the closest vector in $\Lambda^\vee$ to $\boldsymbol{x}$. By using the LWE oracle we can get rid of the errors from approximations and recover $\boldsymbol{x}$, solving the BDD instance.

The second step is quantum, and uses the BDD algorithm above to make a quantum state of a Gaussian $D_{\Lambda,r'}$ with $r' < r/2$. This state can be measured to get a sample from $D_{\Lambda,r'}$. By applying this procedure multiple times we end up with samples which have short length.

5.1 Learning With Errors

In computer science, and in particular in machine learning, there is a problem called *learning from parity with errors* (LPE) which is to find a 'secret' $\boldsymbol{s} \xleftarrow{r} \mathbb{Z}_2^n$ given access to equations

$$\begin{aligned} \langle \boldsymbol{s}, \boldsymbol{a}_1 \rangle &\approx_\varepsilon b_1 \mod 2 \\ \langle \boldsymbol{s}, \boldsymbol{a}_2 \rangle &\approx_\varepsilon b_2 \mod 2 \\ &\vdots \end{aligned}$$

for $\boldsymbol{a}_i \xleftarrow{r} \mathbb{Z}_2^n$ where each equation is correct with probability $1-\varepsilon$. If $\varepsilon = 0$ then this can be solved in polynomial time by, say, Gaussian elimination. To solve this for general $\varepsilon > 0$ we can try the following method: Find a set S of $\boldsymbol{a}_i$-s such that $\sum_S \boldsymbol{a}_i = (1, 0, \ldots, 0)$. Computing the corresponding $\sum_S b_i$ gives a guess for the first bit of $\boldsymbol{s}$. However, this is correct with probability $1/2 + 2^{-O(n)}$. We therefore need $2^{O(n)}$ such guesses to be confident that it is correct. This sketch gives us an algorithm that requires $2^{O(n)}$ equation and runs in $2^{O(n)}$ time. There are improvement to this algorithm, in particular one

by Blum, Kalai and Wasserman [BKW03] who provided a sub-exponential time algorithm which runs in $2^{O(n/\log n)}$ time.

A generalization of LPE is to extend the modulus from 2 to a general integer q. Such a system is described by the following set of equations.

$$\begin{aligned} b_1 &= \langle \boldsymbol{s}, \boldsymbol{a}_1 \rangle + e_1 \\ b_2 &= \langle \boldsymbol{s}, \boldsymbol{a}_2 \rangle + e_2 \\ &\vdots \\ b_m &= \langle \boldsymbol{s}, \boldsymbol{a}_m \rangle + e_m \end{aligned} \tag{6}$$

where $\boldsymbol{a}_i \xleftarrow{r} \mathbb{Z}_q^n$, the $e_i \leftarrow \chi$ for some distribution χ over $\mathbb{R}$ and $\boldsymbol{s} \xleftarrow{r} \mathbb{Z}_q^n$. The addition in the b_i terms is in $\mathbb{R}/q\mathbb{Z}$, i.e. modulo q. The set of equations become non-singular with high probability when m is slightly larger than n. Again, without the error terms this can be solved easily in polynomial time. The problem becomes much harder when we include the error terms.

Definition 5.1 (LWE sample)**.** Pick $\boldsymbol{s} \xleftarrow{r} \mathbb{Z}_q^n$ and $e \leftarrow \chi$ for a distribution χ over $\mathbb{R}$. An *LWE-sample* is a pair (a, b) where $\boldsymbol{a} \xleftarrow{r} \mathbb{Z}_q^n$ and $b = \langle \boldsymbol{s}, \boldsymbol{a} \rangle + e \bmod q\mathbb{Z}$. We say that (a, b) is sampled from the distribution $A_{\boldsymbol{s},\chi}$ over $\mathbb{Z}_q^n \times \mathbb{R}/q\mathbb{Z}$.

We define two problems related to LWE. The first is the recover $\boldsymbol{s}$ which one may think of as the secret key in a scheme based on LWE. The second problem is to distinguish LWE samples from uniform samples.

Definition 5.2 (Search LWE)**.** Given access to m LWE-samples from $A_{\boldsymbol{s},\chi}$, recover $\boldsymbol{s}$.

Definition 5.3 (Decision LWE)**.** Given access to m LWE-samples $(\boldsymbol{a}_i, b_i)$, determine whether they are samples from $A_{\boldsymbol{s},\chi}$ or samples where from the uniform distribution on $\mathbb{Z}_q^n \times \mathbb{R}/q\mathbb{Z}$.

We note that these problems are originally defined so that χ is sampled form a distribution of distributions, but we omit this detail for clarity. If we can show that decision-LWE is hard then LWE samples look uniform to a potential attacker. By doing this we can prove that a cryptographic scheme is semantically secure. We will see an example of this later.

If we assume that the dimension n is not too large and the modulus q is prime and polynomial in n then the search and decision versions of LWE are equivalent.

Proposition 5.4. Assume that the modulus $q = \text{poly}(n)$ and prime. Then search-LWE=decision-LWE.

Proof. search-LWE$\leq$decision-LWE: This is trivial because if we can recover $\boldsymbol{s}$, then we can calculate $\boldsymbol{e} = \boldsymbol{b} - A\boldsymbol{s}$ and check if $\boldsymbol{e}$ is uniformly distributed.

decision-LWE$\leq$search-LWE: Let $\boldsymbol{s} = (s_1, \ldots, s_n)$. Pick an $r \xleftarrow{r} \mathbb{Z}_q$. Transform the sample $(\boldsymbol{a}, b)$ to

$$\begin{aligned} \boldsymbol{a}' &= \boldsymbol{a} + (r, 0, \ldots, 0) \\ b' &= \langle \boldsymbol{s}, \boldsymbol{a} \rangle + e + r \cdot k \end{aligned}$$

where $k \in \mathbb{Z}_q$. Since r was chosen uniformly, $\boldsymbol{a} + (r, 0, \ldots, 0)$ is also uniform. Similarly, $r \cdot k$ will also be uniformly distributed so b' becomes a uniform sample because multiplication by k is a bijection when q is prime. However, if $k = s_1$ we get that

$$\begin{aligned} b_i' &= b + r \cdot k = \langle \boldsymbol{s}, \boldsymbol{a} \rangle + e + r \cdot s_1 \\ &= \langle \boldsymbol{s} + (s_1, 0, \ldots, 0), \boldsymbol{a} + (r, 0, \ldots, 0) \rangle + e \end{aligned}$$

which is another sample from $A_{\boldsymbol{s},\chi}$. We can now use the decision-LWE oracle to determine what the distribution of $(\boldsymbol{a}', b')$ for different values of k. When $k = s_1$ the oracle will decide that the sample is from $A_{\boldsymbol{s},\chi}$ and we have recovered s_1. Since there are only q values k can take and $q = \text{poly}(n)$, this happens fairly quickly. If do a similar step for the other coordinates of $\boldsymbol{s}$ to recover the whole vector. □

This means that if q is prime we can regard both problems as one, which we will denote by LWE. Notice that we always have that decision-LWE$\leq$search-LWE. We also have average case hardness of decision-LWE, which means that if an attacker is able to solve a random instance of LWE, then they can solve *any* instance of LWE.

Proposition 5.5. decision-LWE is as hard in the average case as it is in the worst case.

Proof. Let $\mathcal{A}_{\boldsymbol{s},\chi}$ denote the distribution of an LWE-sample. Assume there is some set $S \subseteq \mathbb{Z}_q^n$ such that LWE-samples where $\boldsymbol{s}' \xleftarrow{r} S$, decision-LWE is easy. Assume also that $|S|/|\mathbb{Z}_q^n| = 1/\text{poly}(n)$, i.e. S is not too small. This

means that we have a distinguisher A which can distinguish samples drawn from $\mathcal{A}_{\boldsymbol{s}',\chi}$ from uniform samples for $\boldsymbol{s}' \xleftarrow{r} S$. It is easy to see that the sample $\{\boldsymbol{a}, \boldsymbol{b} + \langle \boldsymbol{a}, \boldsymbol{t} \rangle / q\}$ is a sample from $\mathcal{A}_{\boldsymbol{s}+\boldsymbol{t}}, \chi$. Therefore, given any $s \xleftarrow{r} \mathbb{Z}_q^n$, pick a uniformly random $\boldsymbol{t} \xleftarrow{r} \mathbb{Z}_q^n$, and query A on

$$\{\boldsymbol{a}, \boldsymbol{b} + \langle \boldsymbol{a}, \boldsymbol{t} \rangle / q\}.$$

Since S is sufficiently large, we will, with high probability, eventually have that $\boldsymbol{s} + \boldsymbol{t} \in S$ and A will be able to decide from which distribution (a, b) is sampled from. This means that if there exists a sufficiently large set S where decision-LWE is easy, then decision-LWE is easy *for all* of $\mathbb{Z}_q^n$. □

We will show later that there exists reductions from lattice problems that we assume are hard to LWE (see Figure 5), meaning that if LWE is easy in the average case, then so is many conjectured hard lattice problems. We conclude that LWE is average-case hard. Note that, by the above proof, there might exists a set S of negligible size where LWE is easy, but as n gets large this set is vanishingly small.

5.1.1 Other Versions of LWE

We mention two other versions of LWE which are also used. Instead of sampling continuous errors we can let χ be a distribution over $\mathbb{Z}_q$, sample $\boldsymbol{a}$ and $\boldsymbol{s}$ like before and let

$$b = \langle \boldsymbol{a}, \boldsymbol{s} \rangle + e \in \mathbb{Z}_q.$$

We call this the *discrete* variant of LWE.

The original definition given in [Reg09] was to let χ be a distribution over $\mathbb{R}/\mathbb{Z}$ (modulo 1) and set an LWE sample to be

$$\begin{aligned} \boldsymbol{a} &\xleftarrow{r} \mathbb{Z}_q^n \\ b = \langle \boldsymbol{a}, \boldsymbol{s} \rangle + e \mod \mathbb{Z} \end{aligned}$$

for $\boldsymbol{s} \xleftarrow{r} \mathbb{Z}_q^n$. We will continue to use the version from Definition 5.1. We do not expect any of these versions to provide different security, and it was shown in [LPR13, Reg09] that all three variants enjoy the same hardness.

5.1.2 Attack

We now give examples of two attacks on LWE which will provide us with some restrictions on the error distribution. Given a LWE-sample $(\boldsymbol{a}, b = \langle \boldsymbol{s}, \boldsymbol{a} \rangle + e)$ and a divisor $q'|q$, we can reduce the sample modulo q' to obtain

$$\begin{aligned} \boldsymbol{a}' &= \boldsymbol{a} \mod q'\mathbb{Z} \\ b' &= \langle \boldsymbol{s}', \boldsymbol{a}' \rangle / q + e \mod q'\mathbb{Z} \end{aligned}$$

where $\boldsymbol{s}' = \boldsymbol{s} \mod q' \cdot \mathbb{Z}$. If we let $q' = 1$ then $\langle \boldsymbol{s}', \boldsymbol{a}' \rangle = 0$ so $b = e \mod \mathbb{Z}$. We now have two potential attacks. Firstly, a distinguishing attack. Because reducing modulo a divisor of q is a map satisfying the condition of Proposition 4.12, $b_i \mod q'$ is uniform if b_i is. Checking whether $b_i \mod q'$ is non-uniform is therefore a distinguishing attack.

Secondly, we have a potential search attack. Assume that the errors from χ usually does not wrap around modulo $\mathbb{Z}$, that is the probability that an error is outside the interval $[-1/2, 1/2)$ is small. Symbolically,

$$\Pr\left[e \notin \left[-\frac{1}{2}, \frac{1}{2}\right)\right] \leq \varepsilon$$

for a small ε. Let $\hat{e}$ be an such an error. Because $\hat{e}$ does not wrap around, $b - \hat{e} = \langle \boldsymbol{s}, \boldsymbol{a} \rangle$ *modulo* q. This gives us an error-less LWE-sample, and collecting enough such errors gives a system of linear equations which can be solved easily. Since the probability that an error does not wrap around is small, we can find sufficiently many such errors quickly.

These two attack also work similarly for other divisors of q that is not too large, but for simplicity we have described it for $q' = 1$. From these two attack we get two requirements for the error distribution χ: It must be statistically indistinguishable from uniform modulo $\mathbb{Z}$, and the sufficiently many errors must wrap around modulo $\mathbb{Z}$.

Let $\chi = D_r$ be the Gaussian of width r exceeding the smoothing parameter of $\mathbb{Z}$, $r \geq \eta_\varepsilon(\mathbb{Z})$. We can for instance let $r = \alpha q > 2\sqrt{n} \geq \eta_{2^{-n}}(\Lambda)$ as described in Theorem 5.6. These errors wrap around modulo $\mathbb{Z}$ with high probability because the distribution is sufficiently wide and $b \mod \mathbb{Z}$ are statistically close to uniform by the choice of the width of D_r (Proposition 4.17).

5.1.3 Hardness

We want to base the hardness of LWE on well known problems for lattices. This can be done by a reduction from LWE, where the errors are sampled

from Gaussian distributions, to SIVP_γ with $\gamma = \mathrm{poly}(n)$. This reduction has a quantum step.

Theorem 5.6. *Let $\varepsilon = \varepsilon(n)$ be some negligible function in n, $\alpha \in (0,1)$ be a real, $q = q(n)$ be some integer such that $q\alpha > 2\sqrt{n}$. Given an efficient (possibly quantum) algorithm that solves LWE_{q,D_r}, then there exists an efficient quantum algorithm for solving $\mathrm{SIVP}_{\tilde{O}(n/\alpha)}$ and $\mathrm{SVP}_{\tilde{O}(n/\alpha)}$.*

By being able to sample from a Gaussian $D_{\Lambda,r}$ of sufficiently small width we can sample a small vector in the lattice Λ. Regev's reduction[Reg09] does this by making use of an LWE oracle. Specifically, given n^c samples from $D_{\Lambda,r}$ we use a LWE oracle to get n^c samples from $D_{\Lambda,r\cdot\sqrt{n}/\alpha q}$. With the condition that $\alpha q > 2\sqrt{n}$ this gives us that

$$r' := r \cdot \frac{\sqrt{n}}{\alpha q} < r \cdot \frac{\sqrt{n}}{2\sqrt{n}} = r/2.$$

The requirement on αq makes sure that we need to perform this iterative step not too many times, as we reduce the width by a factor of 2 each step reducing the width exponentially fast.

Theorem 5.7 (Iterative step [Reg09])**.** *Let $\alpha > 0$ and $q \geq 2$ be an integer. There exists an efficient quantum algorithm that, given a lattice Λ and a number $r > \sqrt{2}q \cdot \eta_\varepsilon(\Lambda)$ for some negligible $\varepsilon = \varepsilon(n)$ such that $r' := r \cdot \omega(\sqrt{\log n})/(\alpha q)$, an oracle to LWE_{q,D_r}, and n^c samples from the discrete Gaussian distribution $D_{\Lambda,r}$, outputs n^c samples from $D_{\Lambda,r'}$.*

An overview of the algorithm is shown in Figure 10. We try here to fill in the details. Essentially, we want to generate samples from a Gaussian distribution $D_{\Lambda,r'}$ given samples from $D_{\Lambda,r}$ of larger width. If we can sample from $D_{\Lambda,r/q}$ we are done, as r/q is even smaller than r'. To do this, draw n^c samples from D_r and divide them by q. However, this gives us samples $\boldsymbol{y} \leftarrow D_{\Lambda/q,r/q}$ which is close, but not quite what we want to sample from. Define the distribution $(\boldsymbol{a}, \boldsymbol{y}) \leftarrow \tilde{D}$ as

$$\begin{aligned} &\boldsymbol{y} \leftarrow D_{\Lambda/q,r/q} \\ &\boldsymbol{a} \xleftarrow{r} \mathbb{Z}_q^n \text{ such that } \boldsymbol{y} \in \Lambda + \Lambda\boldsymbol{a}/q. \end{aligned}$$

By scaling the lattice by $1/q$ we partition it in q^n parts, which means that we can describe the lattice Λ/q by q^n translations of Λ: $\Lambda/q = \{\Lambda + \Lambda\boldsymbol{a}/q\}$ where $\boldsymbol{a} \in \mathbb{Z}_q^n$. An illustration of this is shown in Figure 9.

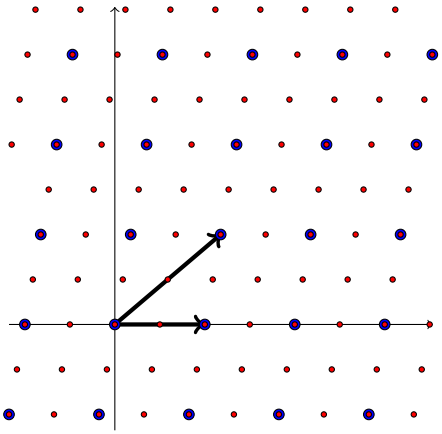

Figure 9: A lattice (large dots) and its scaled counterpart (small dots). Scaling factor $q = 2$ yelds $q^2 = 4$ translates of the lattice.

To obtain the samples $\boldsymbol{y} \leftarrow D_{\Lambda+\Lambda\boldsymbol{a}/q,r/q}$ we just take the original samples from D_r and divide by q. It can be shown that that the distribution of $\boldsymbol{a}$ sampled this way is essentially uniform. We can therefore modify our sampling scheme to

$$\begin{aligned} \boldsymbol{a} &\xleftarrow{r} \mathbb{Z}_q^n \text{ uniformly} \\ \boldsymbol{y} &\leftarrow D_{\Lambda+\Lambda\boldsymbol{a}/q,r/q} \end{aligned}$$

which gives us, essentially, the same distribution. According to [Reg09], a routine calculation shows that the Fourier transform of $D_{\Lambda+\Lambda\boldsymbol{a}/q,r/q}$ is given by

$$\exp(2\pi i\langle \boldsymbol{a}, \tau(\boldsymbol{x})\rangle) \cdot f_{q/r}(\boldsymbol{x}) \tag{7}$$

where

$$\tau(\boldsymbol{x}) = (B^\vee)^{-1}\kappa_{\Lambda^\vee}(\boldsymbol{x}) \mod q$$

that is, the coefficient vector of the closest vector $\kappa_{\Lambda^\vee}(\boldsymbol{x})$ of $\boldsymbol{x}$ in $\Lambda^\vee$. This is the Fourier transform of $D_{\Lambda,r/q}$ multiplied by a phase. Now, recovering $\tau(\boldsymbol{x})$ would mean that we have found the closest vector to $\boldsymbol{x}$ *modulo* q. By using the algorithm described in Figure 11 we can then recover $\boldsymbol{x}$.

To recover $\tau(\boldsymbol{x})$ we use the LWE oracle in the following way. From the definition of the Fourier transform we can view the Fourier transform of $D_{\Lambda+\Lambda\boldsymbol{a}/q}$

as the expectation of $\exp(2\pi i\langle \boldsymbol{x}, \boldsymbol{y}\rangle)$ over $\boldsymbol{y} \leftarrow D_{\Lambda+\Lambda \boldsymbol{a}/q, r/q}$. In symbols

$$\mathrm{E}\left[\exp(2\pi i\langle \boldsymbol{x}, \boldsymbol{y}\rangle)\right] = \exp(2\pi i\langle \boldsymbol{a}, \tau(\boldsymbol{x})\rangle/q) \cdot f_{q/r}(\boldsymbol{x})$$

by (7). Now take $\boldsymbol{x} \in \Lambda^\vee$. The above equation means that

$$\langle \boldsymbol{x}, \boldsymbol{y}\rangle = \langle \boldsymbol{a}, \tau(\boldsymbol{x})\rangle/q \mod 1$$

deterministically, which means that

$$\langle \boldsymbol{a}, \tau(\boldsymbol{x})\rangle = q \cdot \langle \boldsymbol{x}, \boldsymbol{y}\rangle \mod q.$$

Sampling enough $\boldsymbol{y} \leftarrow D_{\Lambda+\Lambda \boldsymbol{a}/q, r/q}$ gives us enough linear equations and we are able to recover $\tau(x)$. Now in the more interesting case that $\boldsymbol{x} \notin \Lambda^\vee$ then we get equations

$$\langle \boldsymbol{a}, \tau(\boldsymbol{x})\rangle \approx \lfloor q \cdot \langle \boldsymbol{x}, \boldsymbol{y}\rangle \rceil \mod q$$

where the rounding is because of the $f_{q/r}$-term in (7). To recover $\tau(\boldsymbol{x})$ we now use the LWE-oracle on these samples. To do this rigorously we need to be assured that $(\boldsymbol{a}, \boldsymbol{y})$ are indeed samples from a valid distribution to make use of the LWE-oracle. See [Reg09] for details.

For completion we describe the second part of the iterative step. To do this we start by creating a quantum state

$$\sum_{\boldsymbol{x} \in \mathbb{R}^n} f_{1/r}|\boldsymbol{x}\rangle$$

corresponding to the Fourier transform of $D_{\Lambda,r}$. We then transform this back into a state corresponding to

$$\sum_{\boldsymbol{y} \in \Lambda} D_{\Lambda,r}(\boldsymbol{y})|\boldsymbol{y}\rangle \tag{8}$$

by using the quantum Fourier transform. However, to do this in a 'reversible' way, we need a BDD-oracle on $\Lambda^\vee$. We can use the oracle described in the first step, and therefore create this state. By measuring the state (8) we can get a sample from $D_{\Lambda,r}$.

Why does this algorithm stop? When we recover $\boldsymbol{x}$, its distance from the dual lattice $\Lambda^\vee$ has to be smaller than $1/r$, this is the condition $r\sqrt{2}q\eta_\varepsilon(\Lambda)$. The iterative steps makes r smaller resulting in larger $1/r$. Eventually $1/r$ is so large that we no longer can solve $\mathrm{BDD}_{\Lambda^\vee, 1/r}$. The algorithm therefore stops.

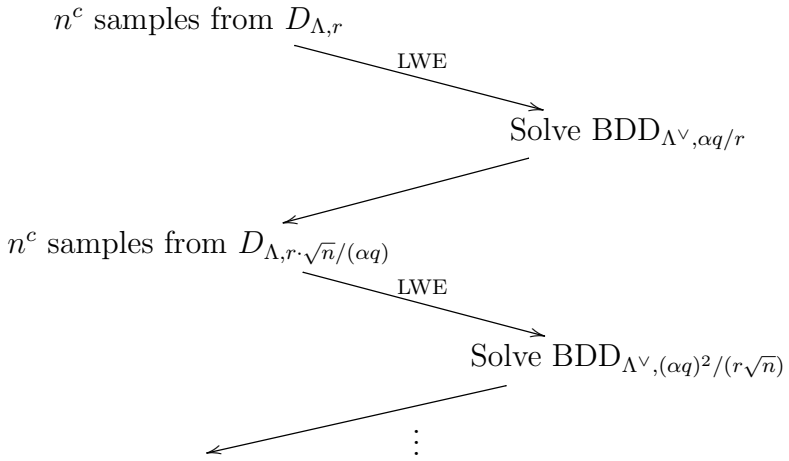

Figure 10: Iterative Step. We produce samples from the discrete Gaussian D_r of progressively narrower width.

5.1.4 More remarks

In the above reduction we showed how to solve BDD on the dual lattice *modulo* q.

Definition 5.8. The q-$\mathrm{BDD}_{\Lambda,d}$ problem is: given an instance $\boldsymbol{y}$ of $\mathrm{BDD}_{\Lambda,d}$ that has solution $\boldsymbol{x}$, find $\boldsymbol{x} \mod q$

It is indeed enough to solve $\mathrm{BDD}_{\Lambda,d}$ modulo q.

Proposition 5.9. There is a polynomial reduction $\mathrm{BDD}_{\Lambda,d}$ from to q-$\mathrm{BDD}_{\Lambda,d}$.

Proof. The algorithm is described in Figure 11. We are given as input a point $\boldsymbol{x} \in \Lambda$ within distance d of Λ. Denote by $\kappa(\boldsymbol{x})$ the closest lattice point to $\boldsymbol{x}$. Now we want to define a sequence $\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots$ in the following way: Let $\boldsymbol{c}_i = B^{-1}\kappa(\boldsymbol{x}_i)$ be the coefficient vector of the the closest vector of $\boldsymbol{x}_i$. Notice that we do not know $\boldsymbol{c}_i$, but we can access a q-BDD oracle to find $\boldsymbol{c}_i \mod q$. Do this, and define

$$\boldsymbol{x}_{i+1} = (\boldsymbol{x}_i - \boldsymbol{c}_i \mod q))/q.$$

Notice that the coefficient-vector

$$\boldsymbol{c}_{i+1} := (\boldsymbol{c}_i - (\boldsymbol{c}_i \mod q))/q \in \mathbb{Z}_q^n$$

is such that $B\boldsymbol{c}_{i+1} \in \Lambda$ and is the closest vector to $\boldsymbol{x}_{i+1}$. Again, we do not know $\boldsymbol{c}_i$, and these are exactly the vectors we want to compute. Now, the distance from $\boldsymbol{x}_{i+1}$ to Λ is at most d/q^i since we assumed that $\boldsymbol{x}_1$ was within d of Λ. After n steps we have found a point $\boldsymbol{x}_{n+1}$ which is within distance d/q^n of Λ. Using a polynomial time algorithm for CVP_γ such as Babai's nearest plane algorithm[Bab86] we can recover the lattice point closest to $\boldsymbol{x}_{n+1}$ with approximation factor 2^n. This yields a lattice point $B\boldsymbol{c}$ within distance

$$2^n \cdot d/p^n < d < \lambda_1(\Lambda)/2$$

of $\boldsymbol{x}_{n+1}$. Hence, $B\boldsymbol{c}_{n+1}$ is the closest point to $\boldsymbol{x}_{n+1}$. Retracing our steps by computing

$$\boldsymbol{c}_n = p\boldsymbol{c}_{n+1} + (\boldsymbol{c}_n \bmod q)$$

gives us $\boldsymbol{c}_n, \boldsymbol{c}_{n-1}, \ldots, \boldsymbol{c}_1$. Since $\boldsymbol{c}_1$ was the coefficient vector of the closest lattice point of $\boldsymbol{x}_1 = \boldsymbol{x}$, $B\boldsymbol{c}_1$ is the closest lattice point to $\boldsymbol{x}$. □

```
1  BDD^{q-BDD}(B, x)
2  x_1 ← x
3  for i = 1 to n do
4  |   c_i mod q ← q-BDD(B, x_i)
5  |   x_{i+1} ← (x_i − B(c_i mod q))/q
6  end
7  c_{n+1} ← CVP_{2^n}(x_{n+1})
8  for n = 1 to 1 do
9  |   c_n ← q c_{n−1} + (c_n mod q)
10 end
11 return Bc_1
```

Figure 11: BDD to q-BDD reduction.

5.1.5 Classical Reduction

The algorithm and reduction above from SIVP_γ to LWE_{q,D_r} is quantum, but we also have a classical polynomial time reduction from a particular lattice

problem problem to LWE. We introduce this for perspective and completion without going into any detail. The algorithm is due to Peikert[Pei09]. We define a special version of the decision version of SVP_γ.

Definition 5.10 (ζ-to-γ-GapSVP)**.** For functions $\zeta(n) \geq \gamma(n) \geq 1$, basis B of lattice Λ and a distance d such that $1 \leq d \leq \zeta(n)/\gamma(n)$, we are asked to determine whether $\lambda_1(\Lambda) \leq d$ or $\lambda_1(\Lambda) > \gamma(n) \cdot \lambda_1(\Lambda)$. We say that it is a YES instance in the former case and a NO instance in the latter.

In this problem we are tasked to determine the length of the shortest vector relative to various parameters, without needing to find it. Obviously there is a trivial reduction from ζ-to-γ-GapSVP to SVP_γ. If we know the shortest vector, we can just calculate the length and decide. Peikert proved a classical reduction from ζ-to-γ-GapSVP to LWE. However because of the introduction of the additional parameter $\zeta(n)$, this reduction is not as convincing as we would like it to be. However, if $q \geq O(2^n)$ then there is a reduction to 'ordinary' GapSVP[Pei09].

5.2 Applications

Regev[Reg09] proposes a simple public-key cryptographic scheme based on LWE. It consists of three algorithms $(\mathcal{K}, \mathcal{E}, \mathcal{D})$ for key-generation, encryption and decryption respectively shown in Figure 12. $\mathcal{K}$ generates the secret key $\boldsymbol{s}$ and the public key $(\boldsymbol{a}, b)$. Encryption $\mathcal{E} : \{0,1\} \to \mathbb{Z}_p^n \times \mathbb{Z}_p$ is only applied to bits. Note that we use the discrete version of LWE here.

The public key has size $O(mn \log p) = \tilde{O}(n^2)$ and the private has size $\tilde{O}(n)$ and encryption increases the size of the message by a factor $O(n)$. If we choose the distribution to be the discrete variant of $\chi = D_\alpha$ then the probability of decryption error is negligible. Additionally, it can be shown that if there exists an attacker who can distinguish between encryptions of 0 and 1, then there exists an attacker on decision-LWE which (assuming the modulus q is prime) gives us a quantum attacker on SIVP_γ.

5.3 Learning With Errors Over Rings

Recall that K is a finite extension of $\mathbb{Q}$. The scheme described above has a key size of $\tilde{O}(n^2)$. However, by having the public key $\{\boldsymbol{a}_i\}$ come from a different ring can reduce the key size to $O(n)$. Before we describe the ring variant of LWE we show where the errors are sampled from. The error

$\mathcal{K}(n, q, m, \chi)$
$\boldsymbol{s} \leftarrow \mathbb{Z}_q^n$
for $i = 1, \dots m$ **do**
  $\boldsymbol{a}_i \leftarrow \mathbb{Z}_q^n$
  $e_i \leftarrow \chi$
  $b_i = \langle \boldsymbol{a}_i, \boldsymbol{s} \rangle + e_i \mod q$
end
return $\boldsymbol{s}, (\boldsymbol{a}_i, b_i)$

$\mathcal{D}(\boldsymbol{s}, (\boldsymbol{a}, b))$
$d = b - \langle \boldsymbol{a}, \boldsymbol{s} \rangle$
if d *closest to* 0 **then**
  return 0
else if d *closest to* $\lfloor \frac{q}{2} \rfloor$ **then**
  return 1
end

$\mathcal{E}(m, (\boldsymbol{a}_i, b_i))$
$S \leftarrow$ subset of $\{1, \dots, m\}$
$\boldsymbol{a} = \sum_{i \in S} \boldsymbol{a}_i$
if $m = 0$ **then**
  $b = \sum_{i \in S} b_i$
else if $m = 1$ **then**
  $b = \lfloor \frac{q}{2} \rfloor + \sum_{i \in S} b_i)$
end
return $(\boldsymbol{a}, b)$

Figure 12: Key generation, encryption and decryption of Regev's scheme [Reg09].

terms are sampled from the space $K_{\mathbb{R}} := K \otimes \mathbb{R}$. By fixing a $\mathbb{Q}$-basis for K, $\{u_1, \dots, u_n\}$, we can write any $\alpha \in K$ as $\alpha = q_1 u_1 + \cdots + q_n u_n$ for $q_i \in \mathbb{Q}$. Now any element $e \in K_{\mathbb{R}}$ is of the form

$$e = (q_1, q_2, \dots, q_n) \otimes x = (q_1 x, q_2 x, \dots, q_n x)$$

for $x \in \mathbb{R}$. Since $q_i x \in \mathbb{R}$, we get that $K_{\mathbb{R}}$ is a vector space over $\mathbb{R}$ with n linearly independent elements and hence isomorphic to $\mathbb{R}^n$ as vector spaces. Intuitively, we want to sample elements s and a from a discrete space and add some 'continuous' noise. Since $K_{\mathbb{R}} \simeq \mathbb{R}^n$ it acts as this continuous domain.

Definition 5.11 (Ring-LWE (R-LWE))**.** Let q be an integer, $\mathcal{I}_q = \mathcal{I}/q\mathcal{I}$ for any ideal $\mathcal{I}$ and χ be a distribution over $K_{\mathbb{R}}$. Pick a secret $s \xleftarrow{r} \mathcal{O}_q^\vee$ and sample $a \xleftarrow{r} \mathcal{O}_q$. An R-LWE sample is the pair (a, b) where

$$b = a \cdot s + e \mod q\mathcal{O}^\vee \in K_{\mathbb{R}}/q\mathcal{O}^\vee.$$

As before, the *search* version is the recover s and the *decision* version is to distinguish an R-LWE sample (a, b) from one where $(a, b) \xleftarrow{r} \mathcal{O}_q \times K_{\mathbb{R}}/q\mathcal{O}^\vee$.

R-LWE is a special case of LWE in the following way: Let B be a $\mathbb{Z}$-basis for $\mathcal{O}$, then for any $a \in \mathcal{O}$, we get that multiplication by a is represented by a matrix relative to the basis B, say A_a, and similarly for s. Now, given a R-LWE sample $(a, b = s \cdot a + e \mod q\mathcal{O}^\vee)$ we get n LWE samples

$$(A_a, \boldsymbol{b} = A_a \cdot \boldsymbol{s} + \boldsymbol{e}).$$

after fixing bases, one for each coordinate. Notice that it is not obvious that the distribution of the columns of A_a are uniformly distributed or that the coordinates in $\boldsymbol{e}$ is independent.

By sampling a_i and s from $\mathcal{O}$ and $\mathcal{O}^\vee$ respectively we have reduced the key size by a factor of n. However, we do not have the same hardness theorems from Section 5.1 because we have introduced more structure. However, [LPR10] shows that we can get similar reductions depending on the parameters and the amount of structure on $\mathcal{O}$. See Section 5.3.3 for further discussion. Notice that the secret s comes from the *dual* ideal $\mathcal{O}^\vee$. We will come back to why this is the 'correct' ring to sample from in Section 5.3.6.

5.3.1 Other Versions of R-LWE

Let ϕ be a distribution over $\mathcal{O}_q^\vee$. We can define a discretized version of R-LWE where we generate samples (a, b) by sampling a and s as usual, sampling $e \leftarrow \phi$ and computing

$$b = a \cdot s + e \in \mathcal{O}_q^\vee.$$

It is clear that if there is an attack on the discrete variant of R-LWE then there is an attack on the variant in Definition 5.11. Any attacker can simply discretize the samples. We might want to use the discrete variant in applications to get exact representation of elements, e.g. an error in $\mathcal{O}^\vee$ can be represented by the integer coefficient of a basis for $\mathcal{O}^\vee$.

We can also define a variant akin to the variant of LWE where $b \in \mathbb{R}/\mathbb{Z}$ is modulo 1. For this version of R-WE the error distribution is over $K_{\mathbb{R}}/\mathcal{O}$. It was shown in [LPR13] that attacks on either of these variants leads to an attack on R-LWE defined in Definition 5.11. We continue using this definition forward.

5.3.2 Error Distribution

In the reduction from SIVP_γ to LWE in the Section 5.1 we saw that we required error samples to be from D_r of appropriate width to guarantee hardness. Since we can view an R-LWE error $e \in K_\mathbb{R}$ as a vector $\boldsymbol{e} \in \mathbb{R}^n$ by fixing bases and get n LWE samples, we want a similar distribution on the coordinates of $\boldsymbol{e}$ We therefore choose the distribution on $K_\mathbb{R}$ such that each coordinate of $\boldsymbol{e}$ is sampled from D_r with similar parameters as the hardness-theorem of LWE. Call this distribution $D_{\boldsymbol{r}}$.

5.3.3 Attack

The simplest attack on R-LWE is to reduce R-LWE samples to regular LWE-samples by fixing bases as described earlier. Each R-LWE sample then gives us n LWE samples. If an attacker can recover useful information from these LWE-samples, such as the secret $\boldsymbol{s} \in \mathbb{Z}_q^n$, then it is easy to recover the corresponding elements in the rings. We therefore need to have that any instantiations of R-LWE cannot be transformed into insecure instantiations of LWE.

Similar to the two attacks on LWE described in Section 5.1 we have two potential attacks on R-LWE, one on search and one on decision. Let $\mathfrak{q}$ be an ideal divisor of $\mathcal{O}^\vee$. Given R-LWE samples

$$(a_i, b_i = s \cdot a_i / + e_i \mod q\mathcal{O}^\vee)$$

we transform them into samples modulo $\mathfrak{q}$ by setting $a_i' = a_i \mod \mathfrak{q}$ and $b_i' = b_i \mod \mathfrak{q}$. Now, $b_i' = s' \cdot a_i' + e_i$ for $s' = s \mod \mathfrak{q}$. We have that reduction modulo $\mathfrak{q}$ satisfies the condition from Proposition 4.12 and therefore maps uniform samples to uniform samples. Now if $\chi \mod \mathfrak{q}$ is detectably non-uniform, we immediately have a distinguishing attack. For each candidate $\hat{s} \in \mathcal{O}/\mathfrak{q}$ for s' check whether $b_i' - \hat{s} \cdot a_i'$ are non-uniform. If such an $\hat{s}$ exists, conclude that the distribution on the original b_i is *not* uniform. If $\mathfrak{q} = \mathcal{O}$ then we only need to check one such representative.

If χ has one or more coordinates that does not 'wrap around' modulo $\mathfrak{q}$, then we can attack search by reducing the R-LWE samples to error-less LWE samples. If we can do this enough times we can solve LWE by, e.g., Gaussian elimination. Similarly to LWE, we choose χ to be a distribution of n-tuples where each coordinate is a Gaussian of width $r_i \leq r$. Denote this distribution by $D_{\boldsymbol{r}}$. By choosing $r \geq \eta_\varepsilon(\mathfrak{q})$ we get that each coordinate wraps

around with high probability and in addition that $b \bmod \mathfrak{q}$ are statistically close to uniform by Proposition 4.17.

5.3.4 Hardness

Analogously to LWE we have a quantum reduction from SIVP to R-LWE. We present the main theorem of [LPR10].

Theorem 5.12. *Let K be an arbitrary number field of degree n, $\alpha \in (0,1)$ and let $q \geq 2$ be such that $q\alpha \geq 2 \cdot \omega(\sqrt{\log n})$. For some negligible $\varepsilon > 0$ there is a probabilistic polynomial time quantum reduction from DGS_γ to R-LWE$_{q,D_\alpha}$, where*

$$\gamma = \max\left\{\eta_\varepsilon(\mathcal{I}) \cdot (\sqrt{2}/\alpha) \cdot \omega(\sqrt{\log n}), \sqrt{2n}/\lambda_1(\mathcal{I}^\vee)\right\}$$

This theorem is essentially the same as the main hardness theorem from LWE, Theorem 5.6, where necessary modifications are done because we are in the R-LWE setting. We use a similar method as for regular LWE: A reduction from $\text{BDD}_{d,\mathcal{I}^\vee}$ to $\text{LWE}_{q,D_\mathbf{r}}$ on the dual lattice given discrete Gaussian samples, and the quantum step to sample from a discrete Gaussian of narrower width. Doing this step many times allows us to sample a short vector in $\mathcal{I}$. The quantum step is more or less identical to that of [Reg09], and we therefore focus on the first reduction.

Similarly to LWE, we only need to find the solution *modulo* q (See Proposition 5.9). Given a q-$\text{BDD}_{\mathcal{I}^\vee,d}$ instance $y = x + e$ with $x \in \mathcal{I}^\vee$ we want to make use of an R-LWE oracle $\mathcal{L}$ and a DGS oracle $\mathcal{D}_{\mathcal{I},r}$ to recover x. Start by computing t such that $t \cdot \mathcal{I}^{-1}$ is coprime to $\langle q \rangle$. Now we use the R-LWE oracle $\mathcal{L}$ as follows: $\mathcal{L}$ will request samples until it is confident that it has a solution. For each of these requests, sample $z \leftarrow D_{\mathcal{I},r}$ and compute the pair (a,b) by

$$\begin{aligned} a &= \theta_t^{-1}(z \bmod q\mathcal{I}) \\ b &= (z \cdot y)/q + e' \bmod \mathcal{O}^\vee. \end{aligned}$$

When $\mathcal{L}$ is confident that it has a solution it will output $s \in \mathcal{O}^\vee$. Finally compute $\theta_t^{-1}(s) \in \mathcal{I}_q^\vee$. To make use of the LWE oracle we need to guarantee that (a,b) is a valid LWE-sample. We do this for a by showing that it is uniform. Since $r \geq \sqrt{2}q \cdot \eta_\varepsilon(\Lambda)$ we have that $D_{\mathcal{I},r}$ is statistically close to uniform. In other words, if each sample is assigned probability β in the

1 $\mathrm{BDD}^{\mathcal{L}}(a, r)$
2 $t \leftarrow t \in \mathfrak{I}$ such that $t \cdot \mathfrak{I}^{-1}$ coprime to $\langle q \rangle$
3 **while** $\mathcal{L}$ requests samples **do**
4 $\quad z \leftarrow D_{\mathcal{I},r}$
5 $\quad e' \leftarrow D_{\alpha/\sqrt{2}}$
6 $\quad a \leftarrow \theta_t^{-1}(z \mod q\mathcal{I})$
7 $\quad b \leftarrow (z \cdot y)/q + e' \mod R^\vee$
8 $\quad (a, b) \to \mathcal{L}$
9 **end**
10 $s \leftarrow \mathcal{L}$
11 **return** $\theta_t^{-1}(s)$

Figure 13: BDD to R-LWE reduction. R-LWE oracle is denoted as $\mathcal{L}$. Line 8 signifies that we provide the LWE oracle $\mathcal{L}$ a sample (a, b) and wait for the response.

uniform distribution, any $z \leftarrow D_{\mathcal{I},d}$ will be assigned probability $\beta \pm \delta$ for some small δ. If this was not the case, then the distance between the two distributions would not be negligible. Now, since θ_t is a bijection and z is essentially uniform, $a = \theta_t^{-1}(z \mod q\mathcal{I})$ is also statistically close to uniform by Proposition 4.12. To finish this proof we need to show that b is sampled from a distribution statistically close to the desired distribution. This proof if quite technical, see [LPR10] for details.

Without going into detail, if the field K is cyclotomic then we can show that there is a reduction from R-LWE to decision-R-LWE. This is done by finding the secret s relative to one ideal factor of $\langle q \rangle$. Because the field is cyclotomic, we can use the field automorphism to then find s relative to *all* ideal factors of $\langle q \rangle$. This enables us to recover s. Details are in [LPR10, Section 5].

5.3.5 Secure Instantiations of R-LWE

A simple attack is to reduce the R-LWE samples to regular LWE samples as we described above. Let $B^\vee$ be a basis for $\mathcal{O}^\vee$. Then its dual is B such that $\sigma(B)^* = \sigma(B^\vee)^{-1}$ where * denotes conjugate-transpose. Because $\boldsymbol{e} \in \mathbb{R}^n$ is the coefficient vector of $e \in K_\mathbb{R}$ relative to the chosen bases we get that

$$\boldsymbol{e} = \sigma(B^\vee)^{-1} \cdot \sigma(e) = \sigma(B)^* \cdot \sigma(e)$$

Since we are considering ideal lattices, the length of elements are bound from below by $\sqrt{n}$ i.e. therefore $||\sigma(b_j)|| \geq \sqrt{n}$ for the basis element $\boldsymbol{b}_j$. If the errors e were sampled from a continuous Gaussian D_r then its counterpart $\boldsymbol{e}$ will be sampled from $\sigma(B)^* \cdot \sigma(D_r)$. This means, since each $\boldsymbol{b}_j \geq \sqrt{n}$, that each coordinate of $\boldsymbol{e}$ is sampled from a Gaussian of width *at least* $r \cdot \sqrt{n}$. By choosing $r \geq 2$ we achieve precisely the hardness condition from LWE[Reg09], meaning that any attack on R-LWE will lead to an attack on LWE. This again gives a quantum algorithm to solve SIVPγ which we have conjectured is hard.

The condition $r \geq 2$ also renders the attack by reducing samples modulo an ideal of $\mathcal{O}^\vee$ useless. The authors of [Pei16] shows that this holds when the norm of the ideal divisor is not too large, specifically less than 2^n which is a pretty mild constrait.

5.3.6 Keys from the Dual Lattice

In our definition of R-LWE, the secret key is an element from the *dual* lattice $\mathcal{O}^\vee$. However, a definition where the samples are all taken from the non-dual $\mathcal{O}$ is equivalent to this definition, up to the error distribution χ[Pei16]. From Proposition 2.23 we have that there exists a bijection

$$\begin{aligned} \theta_t &: \mathcal{O}_q^\vee \to \mathcal{O}_q \\ \theta_t(u) &= t \cdot u, \end{aligned}$$

which we can extend to a map

$$\kappa_t : K_\mathbb{R}/q\mathcal{O}^\vee \to K_\mathbb{R}/q\mathcal{O}$$

naturally. We use κ_t to transform a R-LWE sample $(a_i, b_i = s \cdot a_i + e_i)$ by

$$b_i' = \kappa_t(b_i) = t \cdot b_i = s' \cdot a_i + e_i'$$

and doing nothing with a_i. Let $D_{\boldsymbol{r}}$ be the distribution of e. Notice that $\kappa_t(s) = \theta_t(s)$ because $s \in \mathcal{O}^\vee$, and since θ_t is a bijection, we get that s' is distributed uniformly (because s was). This is therefore a valid R-LWE sample with error distribution $t \cdot D_{\boldsymbol{r}}$ and with secret from the non-dual lattice. Because θ_t is an *efficiently invertible* bijection, finding s' immediately yields s, so solving search for the transformed sample (with errors from $t \cdot D_{\boldsymbol{r}}$) is equivalent to solving search for the original samples. Additionally, because κ_t

sends uniform samples to uniform samples we get that the decision versions are equivalent as well.

However, even though we are able to transform a 'dual' sample to a 'non-dual' sample, the errors come from a new error distribution $t \cdot \chi$. So if $D_{\boldsymbol{r}}$ itself satisfies the hardness properties of R-LWE, $t \cdot D_{\boldsymbol{r}}$ might not. Consider an error $e \leftarrow D_{\boldsymbol{r}}$. Under κ_t this becomes $t \cdot e$ and has norm

$$\begin{aligned} ||te|| = ||\sigma(te)|| &= ||(\sigma_1(te), \dots, \sigma_n(te)|| \\ &= || \left(\sigma_1(t)\sigma_1(e), \dots, \sigma_n(t)\sigma_n(e)\right) || \end{aligned}$$

In the trivial case where $t \in \mathbb{Z}$ we get that $t \cdot \chi$ is just a scaled version of χ, meaning a spherical distribution remains spherical but with different width. In general, however, $t \in \mathcal{O}^\vee$. Now if χ is spherical, $t \cdot \chi$ need not be since each i-th coordinate is scaled by $\sigma_i(t)$. We might try to scale each coordinate by the largest $\sigma_i(t)$ such that we get a wider spherical distribution. However, this means that we need to change other parameters of the system to guarantee security or that the errors get so large that we end up with decryption errors.

Recall that if we sample from a Gaussian with width r which exceeds the smoothing parameter for the lattice, decryption becomes impossible because samples are essentially uniform. The amount of noise we can allow while still doing this depends inversely on λ_n of the dual ideal, because

$$\eta_\varepsilon(\Lambda) \leq \sqrt{n}/\lambda_1(\Lambda^\vee) \leq \sqrt{n}/\lambda_n(\Lambda^\vee)$$

from Proposition 4.18. Therefore, by guaranteeing a small bound for $\lambda_n(\Lambda^\vee)$ we get a large smoothing parameter which allows us to sample from a Gaussian with large width. If we restrict ourselves to cyclotomic fields, $\mathcal{O}$ contains n roots of unity $\{1, \zeta, \dots, \zeta^{n-1}\}$ all of norm $\sqrt{n}$. We therefore get that $\lambda_n(\mathcal{O}) = \sqrt{n}$. In addition we have from Proposition 3.14

$$\lambda_1(\mathcal{O}) \geq \sqrt{n}$$

and in particular

$$\lambda_n(\mathcal{O}) \geq \sqrt{n}$$

which means that this is the optimal value for λ_n. Sampling from $\mathcal{O}^\vee$ therefore gives us the largest bound on the smoothing parameter, allowing us to decrypt correctly for large errors.

5.3.7 Applications

We describe a simple cryptographic scheme based on R-LWE and prove that it is semantically secure. Let χ be a distribution over $K_{\mathbb{R}}$. Fix the ring $R = \mathbb{Z}[X]/\langle X^n - 1\rangle$, which is self-dual for n a power of 2. Let $R_q = R/qR$ and $s, e \leftarrow \chi$. We generate keys by choosing $a \xleftarrow{r} R_q$ and setting

$$\begin{aligned} s &\quad \text{as secret key} \\ (a, s \cdot a + e) &\quad \text{as public key.} \end{aligned}$$

The encryption algorithm samples $r, e_1, e_2 \leftarrow \chi$, computes

$$u = a \cdot r + e_1 \mod q \quad \text{and} \quad v = b \cdot r + e_2 + \lfloor q/2 \rceil \cdot z \mod q$$

and outputs (u, v) by viewing a cipher text $z \in \{0,1\}^n$ as coefficients in a polynomial in R. The decryption algorithm, on input (u, v), computes

$$v - u \cdot s = (r \cdot e - s \cdot e_1 + e_2) + \lfloor q/2 \rceil \cdot z \mod q.$$

This is essentially the same scheme as for regular LWE described in Section 5.1. If we chose parameters correctly, the coefficients of $r \cdot e - s \cdot e_1 + e_2$ have magnitude less than $q/4$ (with high probability), so we can recover the bits of z by checking if the coefficient is closer to 0 or to $\lfloor q/2 \rceil$.

It can be show that decision-R-LWE is hard even when $s \leftarrow \chi$. An attacker can therefore not determine whether a pair $(a, b) \in R_q^2$ is from the distribution above of uniform. Therefore, the message an attacker sees is $a \cdot r + e_1 \mod q$ and $b \cdot r + e_2 + \lfloor q/2 \rceil$, which are R-LWE samples with secret r and are also hard to distinguish from uniform samples. This gives us semantic security.

6 Conclusion

In these thesis we have looked various lattice problems and how they relate to each other. We showed simple reductions between standard lattice problems and how they can be attacked to get an overview of their hardness. Then we showed that both LWE and R-LWE are at least as hard as quantum solving hard lattice problems. This means that cryptographic schemes based on LWE and R-LWE are provably hard when instantiated correctly. We also saw how sampling the secret s from the dual ideal $\mathcal{O}^\vee$ in R-LWE is the correct way to instantiate this scheme.

References

[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, pages 99–108, New York, NY, USA, 1996. ACM.

[Ajt98] Miklós 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, pages 10–19, New York, NY, USA, 1998. ACM.

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

[BGV14] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. *ACM Trans. Comput. Theory*, 6(3):13:1–13:36, July 2014.

[Bia15] Jean-François Biasse. A fast algorithm for finding a short generator of a principal ideal of $\{Q\}(\zeta_{\{}p^{s}\})$. 2015.

[BKW03] Avrim Blum, Adam Kalai, and Hal Wasserman. Noise-tolerant learning, the parity problem, and the statistical query model. *J. ACM*, 50(4):506–519, July 2003.

[BS16] Jean-François Biasse and Fang Song. A polynomial time quantum algorithm for computing class groups and solving the principal ideal problem in arbitrary degree number fields. 2016.

[Cai03] Jin-Yi Cai. A new transference theorem in the geometry of numbers and new bounds for ajtai's connection factor. *Discrete Applied Mathematics*, 126(1):9 – 31, 2003. 5th Annual International Computing and combinatorics Conference.

[CDPR16] Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology – EUROCRYPT 2016*, pages 559–585, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg.

[Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing*, STOC '09, pages 169–178, New York, NY, USA, 2009. ACM.

[GGH11] Oded Goldreich, Shafi Goldwasser, and Shai Halevi. *Collision-Free Hashing from Lattice Problems*, pages 30–39. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.

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

[KLWLL82] Arjen K. Lenstra, H W. Lenstra, and László Lovász. Factoring polynomials with rational coefficients. 261, 12 1982.

[LATV17] Adriana López-Alt, Eran Tromer, and Vinod Vaikuntanathan. Multikey fully homomorphic encryption and applications. *SIAM Journal on Computing*, 46(6):1827–1892, 2017.

[LM06] Vadim Lyubashevsky and Daniele Micciancio. Generalized compact knapsacks are collision resistant. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *Automata, Languages and Programming*, pages 144–155, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.

[LMPR08a] Vadim Lyubashevsky, Daniele Micciancio, Chris Peikert, and Alon Rosen. Swifft: A modest proposal for fft hashing. In Kaisa Nyberg, editor, *Fast Software Encryption*, pages 54–72, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.

[LMPR08b] Vadim Lyubashevsky, Daniele Micciancio, Chris Peikert, and Alon Rosen. Swifft: A modest proposal for fft hashing. In Kaisa Nyberg, editor, *Fast Software Encryption*, pages 54–72, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.

[LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded 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.

[LPR13] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for ring-lwe cryptography. Cryptology ePrint Archive, Report 2013/293, 2013. https://eprint.iacr.org/2013/293.

[Ogg10] Frederique Oggier. Introduction to algebraic number theory, 2010. Lecture notes given at NTU 2009-2010, http://www1.spms.ntu.edu.sg/~frederique/ANT10.pdf.

[Pei09] Chris Peikert. Public-key cryptosystems from the worst-case shortest vector problem: Extended abstract. In *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing*, STOC '09, pages 333–342, New York, NY, USA, 2009. ACM.

[Pei16] Chris Peikert. How (not) to instantiate ring-lwe. In Vassilis Zikas and Roberto De Prisco, editors, *Security and Cryptography for Networks*, pages 411–430, Cham, 2016. Springer International Publishing.

[Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):34:1–34:40, September 2009.

[RM07] Oded Regev and D. Micciancio. Worst-case to average-case reductions based on gaussian measures. *SIAM Journal of Compututing*, pages 267–302, 2007.

[RS10] Markus Rückert and Michael Schneider. Estimating the security of lattice-based cryptosystems. 2010:137, 01 2010.

[Sho97] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, October 1997.

[vDGHV10] Marten van Dijk, Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. Fully homomorphic encryption over the integers. In Henri Gilbert, editor, *Advances in Cryptology – EUROCRYPT 2010*, pages 24–43, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.

[Was82] Lawrence C. Washington. *Introduction to Cyclotomic Fields*, volume 83. Springer-Verlag, 1 edition, 1982.