SoliloQuat: Throwing Caution to the Wind

Andrew Mendelsohn and Ben Nelson

Department of EEE, Imperial College London, United Kingdom.
andrew.mendelsohn18@imperial.ac.uk, bennelson3141@gmail.com

Abstract. We propose a plausibly post-quantum additively homomorphic PKE scheme, SoliloQuat, based on the short generator principal ideal problem (SG-PIP) in orders of quaternion algebras. SoliloQuat is inspired by Soliloquy, a KEM that was both introduced and broken by Campbell-Groves-Shepherd in 2014. However, it is not known if their attack can be generalised to the non-commutative setting, despite having received cryptanalytic attention due to a reduction from the rank 2 module-LIP instances underlying HAWK to nrd-PIP (Eurocrypt '25). Demonstrating the correctness of our scheme requires novel results on the eigenvalues of the left regular representation of quaternions, which may be of independent interest. We prove IND-CPA security of our scheme, assuming the hardness both of SG-PIP in orders of quaternion algebras, and some less-exotic lattice-based assumptions.

Keywords: post-quantum cryptography, principal ideal problem, quaternions, lattices

1 Introduction

In Gentry's thesis [14], the author introduced a novel public key encryption (PKE) scheme based on ideal lattices derived from algebraic number fields. The motivation of this scheme was to develop fully homomorphic encryption (FHE),[1] which was achieved using the arithmetic properties of number fields.

Inspired by Gentry's work, Smart and Vercauteren built an FHE scheme [26] which boasted high levels of efficiency and compactness, once key pairs had been generated. Their scheme is a specialisation of Gentry's to the setting of principal ideals, and its security was based on the 'short generator principal ideal problem' (SG-PIP)[2] in cyclotomic fields: in particular, SG-PIP requires one to find a 'short' element of an ideal lattice which generates the ideal (in the ring-theoretic sense). Unfortunately, key generation was a computational bottleneck which became intractable even for relatively small parameters.

[1] FHE is a cryptographic primitive which permits unbounded computation on encrypted data by performing computations on the corresponding ciphertexts. The concept of FHE was introduced to the literature in 1978 [24].

[2] There were prior attempts to build cryptography from PIP, e.g. [3].

Smart and Vercauteren's scheme was subsequently broken. In [4], researchers from GCHQ published a similar proprietary scheme named 'Soliloquy' (in fact a KEM), developed some years previously, together with a sketch of a quantum attack solving SG-PIP in certain cyclotomic fields. This was followed by [7,2], which made the details of the attack more precise, spawning a series of follow-up works solving similar computational problems on ideal lattices (e.g. [8,22,9,1]). Both Smart-Vercauteren and SOLILOQUY thus suffered a complete break.

In the meantime, the lattice isomorphism problem (LIP) has emerged as a promising candidate for post-quantum cryptography. LIP requires one to decide if two lattices, represented by (pseudo)bases, are geometrically equivalent, and if so, to find an isometry between them. To improve the compactness and efficiency of LIP-based schemes, structure has been introduced in the form of Hermitian lattices corresponding to finite rank modules over the ring of integers of cyclotomic fields [10]. It was subsequently shown that these structured LIP instances reduce in polynomial time to a quaternion-ideal problem known as nrd-PIP [6], which is a PIP variant where one must find a generator with prescribed reduced norm. Such a reduction implies that if nrd-PIP in an order of a quaternion algebra is efficiently solvable, then so is rank two module-LIP over the ring of integers of a cyclotomic field. In particular, a solution to nrd-PIP would break HAWK.

Prior work by Lenstra-Silverberg [20] (and Kirchner [17]) gave a deterministic polynomial-time algorithm to solve nrd-PIP for ideal lattices in orders of CM-fields, which includes nrd-PIP for ideals in orders of cyclotomic fields. This algorithm built on work by Gentry and Szydlo [13], rendering nrd-PIP in cyclotomic fields completely unsuitable for any cryptographic application. Notably, however, there is no such efficient algorithm known for the quaternionic analogue of nrd-PIP.

To date, the most efficient algorithms to solve either SG-PIP or nrd-PIP for orders of quaternion algebras require the use of an (approximate) SVP oracle in the corresponding ideal lattice [18]. For quaternionic ideal lattices of sufficiently large rank, solving SG-PIP or nrd-PIP appears intractable. This statement assumes that neither the techniques of Lenstra-Silverberg-Kirchner nor [7,2] can be extended to the relevant quaternion algebras, which are currently open problems. There is a strong obstacle to the existing techniques for solving PIP in number fields straightforwardly lifting to solve quaternionic analogues of PIP: in particular, the unit group of a quaternion order is noncommutative. This appears to preclude any useful definition of the corresponding log-unit lattice, which is a central tool used in the algorithms of [7,2]. Moreover, the work of Lenstra-Silverberg requires the subgroup of units having reduced norm one to form an abelian group, but this is known to not be the case for the quaternion orders considered below [29].

The prospect of rebuilding provably secure SOLILOQUY-style schemes from quaternion algebras, as opposed to number fields, becomes, very tentatively, a basis for efficient PKE, should the key generation problem of Smart-Vercauteren be overcome. However, such a scheme would have large and obvious candidate attack vectors: in order to break a revived SOLILOQUY variant in quaternion algebras, one could either upgrade Lenstra-Silverberg-Kirchner's work, extend the methods of [7,2], or find novel algorithms for quaternion PIP. Developing such a scheme would be diving headlong into cryptanalytic cut-and-thrust, presenting a sitting duck of a target, under siege from multiple active research directions at once.[3]

Our Contributions In this work we present SOLILOQUAT, a PKE scheme based on PIP in certain orders of quaternion algebras. Our scheme is additively homomorphic, but moving to the setting of quaternions appears to preclude the possibility of obtaining a multiplicatively homomorphic scheme in the absence of new techniques. Our key generation, encryption, and decryption algorithms are loosely analogous to the corresponding algorithms of Smart-Vercauteren, but rigorously proving correctness in the quaternionic setting introduces various technical hurdles to be overcome.

In particular, we require a classification of maximal right ideals in a certain non-maximal order (sometimes referred to as the 'natural order') which has not appeared in the literature before, so that key generation is well defined. We also conduct a close study on the eigenvalue distribution of the left regular representation of quaternions with integer coefficients sampled from a discrete Gaussian distribution, allowing us to obtain bounds on the radius of decryption. This in turn implies correctness of our scheme for well-chosen parameters.

Finally, we demonstrate the security of SOLILOQUAT under worst-case lattice assumptions. We are able to ground the one-wayness of encryption on the approximate bounded distance decoding problem (γ-BDD), while IND-CPA security is based on the decision bounded distance decoding problem (decision BDD) as studied by Gentry [14, Chapter 7.7]. Key recovery directly corresponds to an SG-PIP instance.

Technical Overview For expository purposes, we give a brief and simplified sketch of our scheme. We first introduce quaternion algebras over algebraic number fields; for more background information on these objects, see Section 2.

Let $K \subset \mathbb{R}$ be an algebraic number field. Define quaternionic basis elements by the relations $i^2 = j^2 = -1$ and $ij = -ji = k$. We then set

$$\mathcal{A} = K + iK + jK + kK,$$

[3] To mix metaphors.

and denote $n := [\mathcal{A} : \mathbb{Q}]$. Naturally, $\mathcal{A}$ contains the quadratic K-extension $L := K(i)$, and therefore also contains

$$\Lambda = \mathcal{O}_L + j\mathcal{O}_L.$$

The ring (in fact, an order) Λ will be the backdrop for SoliloQuat.

Our key generation algorithm produces keys as follows: sample $\alpha \leftarrow \mathcal{D}_{\Lambda,\sigma}$ from a discrete Gaussian distribution over Λ, until α satisfies $N_{K/\mathbb{Q}}(\mathrm{nrd}(\alpha)) = p$ for some prime $p \equiv 1 \bmod n$, and $2^{(p-1)/n}$ has order n modulo p.[4] Such an element α is the private key in SoliloQuat.

By our classification of maximal right ideals, α generates a maximal right ideal $\alpha\Lambda$ with an explicit form; this classification is given by Proposition 1. In particular, when L is a cyclotomic field and K is the maximal totally real subfield of L, such ideals have a three-element representation (p, c, θ) where $(p, \zeta_n - c)\mathcal{O}_L$ is a prime ideal, and $\theta \in \Lambda$. For the rest of this introduction, and indeed the rest of the paper, we work in the setting $L = \mathbb{Q}(\zeta_n)$, $K = \mathbb{Q}(\zeta_n) \cap \mathbb{R}$, and $n = 2^l$ for some $l \geq 2$. The three-element representation (p, c, θ) forms the SoliloQuat public key.

Suppose that we wish to encrypt a message $\varepsilon \in \{0,1\}^n$. In order to enter the quaternion setting, we associate ε with the element of Λ with coefficients given by the entries of ε. We then generate a short element $\delta \leftarrow 2\Lambda$, defined as a sample from a discrete Gaussian distribution over Λ, multiplied by two. We can then form the element $\varepsilon' := \varepsilon + \delta$, and we output the ciphertext

$$z := \varepsilon' \bmod \alpha\Lambda.$$

One can view z as an element of $\mathbb{F}_p^2$. This element can be efficiently computed using the public three-element representation of $\alpha\Lambda$, and can therefore be computed by anyone with knowledge of the public key.

Decryption works as follows: given a ciphertext z encrypting $\varepsilon \in \Lambda$ and knowledge of the private key α, one computes

$$\varepsilon' = z - \alpha\lfloor\alpha^{-1}z\rceil, \tag{1}$$

where $\lfloor\cdot\rceil$ denotes coordinate-wise rounding.[5] Indeed, by construction of SoliloQuat encryption, there exists $x \in \Lambda$ satisfying

$$\varepsilon' - z = \alpha x. \tag{2}$$

If $\lfloor\alpha^{-1}\varepsilon'\rceil = 0$, then rearranging equation (2) by multiplying both sides by α^{-1} and applying coordinate-wise rounding, we find that $x = -\lfloor\alpha^{-1}z\rceil$. Plugging this

[4] We also check that, in a certain subfield of $\mathcal{A}$, the ideals lying over p are not principal.

[5] Of course, one must specify a basis of Λ in order for this to be well-defined. Further details are given in Section 5.

back into equation (2) yields the decryption formula (1). The subtlety of this work then lies in choosing parameters such that $\lfloor \alpha^{-1}\varepsilon' \rceil = 0$ with high probability (to guarantee correct decryption). This turns out to be rather non-trivial, and consumes both Section 6 and Appendices A.3, A.4 of the paper.

Once we have demonstrated that parameters may be chosen to ensure SoliloQuat ciphertexts may be correctly decrypted, we give an analysis of security. We argue for one-wayness, IND-CPA security, and the hardness of key recovery. We claim one-wayness under $O(n^{5/4})$-BDD. In a similar fashion, IND-CPA security is demonstrated assuming decision BDD. Key recovery is the problem of computing α given the three-element representation (p, c, θ). This is manifestly an SG-PIP problem, since (p, c, θ) generates a principal right Λ-ideal, and α has bounded coefficients.

Discussion In this work we revive the ideas behind [26,4] in the quaternionic setting, giving a PKE scheme that is plausibly post-quantum secure. Security is based on approx-BDD and decision BDD, as well as the somewhat less standard SG-PIP problem for maximal right ideals of Λ. We concede that we do not have strong evidence that the latter problem is hard, and as mentioned above, there are broad and leafy avenues down which a cryptanalyst might wander to arrive at a break. However, given that PIP (in the form of nrd-PIP, not SG-PIP) has re-emerged in lattice-based cryptography, and that the security of HAWK relies on the hardness of this problem, we consider it worthwhile to construct other, efficient, highly functional schemes from plausibly secure instances of PIP. At the very least, we hope our work will encourage further cryptanalysis of variants of quaternion PIP to improve our understanding of the security of HAWK, and also, if we should be so lucky, to bolster our confidence in SoliloQuat.

Paper Organisation In Section 2 we detail the mathematical preliminaries. We then recall Soliloquy in Section 3. In order to give an outline of SoliloQuat in Section 5, we first classify the maximal right ideals of Λ in Section 4. Correctness and security of SoliloQuat occupy Sections 6 and 7. We then conclude, and provide further mathematical detail in the Appendix.

2 Preliminaries

2.1 Notation

Throughout the paper, we adopt the following notation:

- '$\hookrightarrow$' denotes an injective map;
- '$\subset$' denotes set-theoretic inclusion, which need *not* be strict;
- '$M_n(F)$' denotes the $n \times n$ matrix ring over the field F;
- 'I_n' denotes the $n \times n$ identity matrix, regardless of the ambient matrix ring;

- '$\mathbf{e}_r$' denotes the r-th column of I_n whenever the dimension is understood;
- $\lfloor \cdot \rceil$ denotes coordinate-wise rounding, whenever the basis is clear;
- for $x \in \mathbb{C}$, we denote the complex conjugate of x by $\overline{x}$, and the Euclidean norm of x by $|x|$;
- for a prime-power p^r, we denote the unique field of order $q = p^r$ by $\mathbb{F}_q$;
- the algebraic closure of $\mathbb{F}_p$ will be written as $\overline{\mathbb{F}}_p$;
- $\mathrm{neg}(n)$ denotes any function which decays faster than $1/n^c$ for any $c > 0$.

2.2 Lattices

A *lattice* is a discrete additive subgroup of a finite-dimensional vector space V over a field F. By choosing a basis of V, we may assume that $V = F^n$ where $n = \dim_F(V)$. Any lattice $\mathcal{L} \subset F^n$ can be generated as a $\mathbb{Z}$-module by a set of linearly independent vectors $B = \{\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_m\}$, which is known as a *basis*. Given a basis B of $\mathcal{L}$, we can form the associated *generating matrix*

$$\mathbf{B} = \begin{pmatrix} \uparrow & \uparrow & & \uparrow \\ \mathbf{b}_1 & \mathbf{b}_2 & \cdots & \mathbf{b}_m \\ \downarrow & \downarrow & & \downarrow \end{pmatrix}.$$

The generating matrix $\mathbf{B}$ is so-named because

$$\mathcal{L} = \mathbf{B}\mathbb{Z}^m,$$

where $\mathbb{Z}^m$ represents the set of $m \times 1$ vectors with integer entries. We will now discuss the familiar setting of the *geometry of numbers*, where we assume that $F \subset \mathbb{R}$. In this case, the *volume* of $\mathcal{L}$ is defined as

$$\mathrm{Vol}(\mathcal{L}) = |\det(\mathbf{B})|,$$

which doesn't depend on the choice of generating matrix $\mathbf{B}$. Given a vector $\mathbf{x} = (x_r) \in F^n$, we define the 2-norm $||\cdot||_2$ via

$$||\mathbf{x}||_2^2 = \sum_{r=1}^{n} x_r^2.$$

This allows us to define the *successive minima* of $\mathcal{L}$:

Definition 1 (Successive Minima). *Let $\mathcal{L} \subset \mathbb{R}^n$ be a lattice. Given $\lambda > 0$, let $\mathcal{B}_2(\lambda)$ denote the 2-ball in $\mathbb{R}^n$ with radius λ. For $1 \leq r \leq n$, we define*

$$\lambda_r(\mathcal{L}) = \inf\left\{\lambda > 0 : \mathcal{L} \cap \mathcal{B}_2(\lambda) \text{ contains } r \text{ linearly independent vectors}\right\}.$$

Minkowski's theorems are central to both the geometry of numbers and lattice-based cryptography. While these theorems come in many different forms, we are particularly interested in the following variant:

Theorem 1 (Minkowski). *Suppose that $\mathcal{L} \subset \mathbb{R}^n$ is a lattice and let V_n denote the volume of $\mathcal{B}_2(1)$, the unit n-ball in $\mathbb{R}^n$. Then*

$$\frac{2^n}{n!} \operatorname{Vol}(\mathcal{L}) \leq V_n \lambda_1 \lambda_2 \cdots \lambda_n \leq 2^n \operatorname{Vol}(\mathcal{L}). \tag{3}$$

Proof. For example, see [5, pp. 203]. □

We will often sample elements from lattices, which are intrinsically discrete structures. As such, we require a discrete analogue of the Gaussian distribution:

Definition 2 (Discrete Gaussian on $\mathbb{Z}^n$). *Given $\sigma > 0$, define the function $\rho_\sigma : \mathbb{R} \to \mathbb{R}$ by the formula*

$$\rho_\sigma(x) = exp\left(-\frac{x^2}{\sigma^2}\right).$$

Suppose that an integer-valued random variable X has distribution

$$\mathbb{P}(X = r) = \frac{\rho_\sigma(r)}{\sum_{s \in \mathbb{Z}} \rho_\sigma(s)}.$$

We say that X has a discrete Gaussian distribution of mean 0 and width σ, which is denoted by $X \sim \mathcal{D}_{\mathbb{Z},\sigma}$. If $X_r \sim \mathcal{D}_{\mathbb{Z},\sigma}$ for $1 \leq r \leq n$, then we can construct the random variable

$$X := \left(X_1\ X_2 \cdots X_n\right)^\intercal,$$

and we write $X \sim \mathcal{D}_{\mathbb{Z}^n,\sigma}$.

Lattice-based cryptography is one of the most promising avenues for post-quantum cryptography, providing a rich supply of computationally hard problems; we now define some of these worst-case lattice problems:

Definition 3 (γ-BDD). *Let $0 < \gamma = \gamma(n) < 1/2$. Given a lattice $\mathcal{L} \subset \mathbb{R}^n$, a radius $0 < r < \gamma \cdot \lambda_1(\mathcal{L})$, and a vector $\mathbf{t} = \mathbf{v} + \mathbf{e}$ where $||\mathbf{e}||_2 < r$ and $\mathbf{v} \in \mathcal{L}$, output the lattice vector $\mathbf{v}$.*

One closely-related problem to γ-BDD is the *decision bounded distance decoding problem (decision BDD)*, which was studied by Gentry [14, Section 7.7] as a special case of his so-called *ideal coset problem*:

Definition 4 (Decision BDD). *Let $\mathcal{L} \subset \mathbb{Z}^n$ be a lattice, and fix some radius $0 < r \ll \lambda_1(\mathcal{L})/\sqrt{n}$. The challenger samples a random bit $b \leftarrow_R \{0, 1\}$. If $b = 0$, the challenger samples $\mathbf{v} \leftarrow \mathcal{D}_{\mathbb{Z}^n,r}$ and outputs*

$$\mathbf{r} := \mathbf{v} \pmod{\mathcal{L}}.$$

If $b = 1$, the challenger outputs a uniformly random coset $\mathbf{r}$ of the quotient $\mathbb{Z}^n/\mathcal{L}$. Given $\mathbf{r}$, the decision BDD problem is to guess whether $b = 0$ or $b = 1$.

2.3 Number Fields

An *algebraic number field* is a finite extension of the rational numbers $\mathbb{Q}$. If F is an algebraic number field of degree $d = [F : \mathbb{Q}]$, then there are exactly d embeddings $F \hookrightarrow \mathbb{C}$. An embedding $\iota : F \hookrightarrow \mathbb{C}$ is said to be *real* if the image of ι is contained in $\mathbb{R}$; otherwise, ι is said to be *complex*. If all of the embeddings $F \hookrightarrow \mathbb{C}$ are real, then we say that F is *totally real*. Any number field F contains a *maximal totally real subfield* F^+, which is the unique totally real subfield of F that is not properly contained in any other totally real subfield.

Given a pair of number fields $E \subset F$, we define $\mathrm{Aut}(F/E)$ to be the group of bijective ring homomorphisms $F \to F$ that fix E pointwise. If

$$\#\mathrm{Aut}(F/E) = [F : E],$$

then the extension F/E is said to be *Galois*, and we write '$\mathrm{Gal}(F/E)$' in the place of '$\mathrm{Aut}(F/E)$'. Suppose that F/E is a Galois extension, and abbreviate $G := \mathrm{Gal}(F/E)$. In this setting, we may define the *algebraic norm* $N_{F/E} : F \to E$ by the formula

$$N_{F/E}(x) = \prod_{\sigma \in G} \sigma(x),$$

where $x \in F$. For any number field F, an element $x \in F$ is said to be *integral* if x is the root of some monic polynomial $f \in \mathbb{Z}[X]$. The *ring of integers* of F, written $\mathcal{O}_F$, consists of all the integral elements of F, which naturally forms a ring. A finite collection of elements $\{x_r\}_{r \in S}$ forms an *integral basis* if every element $x \in \mathcal{O}_F$ has a unique expression of the form

$$x = \sum_{r \in S} a_r x_r,$$

where the coefficients a_r are rational integers. Every number field F has a (non-unique) integral basis.

We are particularly interested in a class of number fields known as the *cyclotomic number fields*, and their maximal totally real subfields. A cyclotomic number field is formed by adjoining a primitive n-th root of unity to $\mathbb{Q}$, which we will usually denote by $L = \mathbb{Q}(\zeta_n)$. Note that $L/\mathbb{Q}$ has degree $\phi(n)$, where ϕ denotes Euler's totient function. The maximal totally real subfield of L is $\mathbb{Q}(\zeta_n + \zeta_n^{-1})$, which we will usually denote by K. Both L and K are Galois over $\mathbb{Q}$, with respective Galois groups (up to isomorphism)

$$\mathrm{Gal}(L/\mathbb{Q}) \cong (\mathbb{Z}/n\mathbb{Z})^\times,$$
$$\mathrm{Gal}(K/\mathbb{Q}) \cong (\mathbb{Z}/n\mathbb{Z})^\times/\{\pm 1\}.$$

Both of the above isomorphisms can be made explicit. For example, $\mathrm{Gal}(L/\mathbb{Q})$ consists of the automorphisms of L sending $\zeta_n \mapsto \zeta_n^r$ where $(r, n) = 1$. Since ζ_n has order n, we can identify r with an element of $(\mathbb{Z}/n\mathbb{Z})^\times$, which yields the

above isomorphism. An analogous statement can be made for $\mathrm{Gal}(K/\mathbb{Q})$. The rings of integers $\mathcal{O}_L$ and $\mathcal{O}_K$ are precisely

$$\mathcal{O}_L = \mathbb{Z}[\zeta_n],$$
$$\mathcal{O}_K = \mathbb{Z}[\zeta_n + \zeta_n^{-1}].$$

As such, we take the natural integral basis for $\mathcal{O}_L$ to be the *power basis*

$$\mathcal{B} := \{1, \zeta_n, \zeta_n^2, \ldots, \zeta_n^{\phi(n)-1}\}.$$

We also note that $\{1, \zeta_n\}$ is a relative integral basis for $\mathcal{O}_L/\mathcal{O}_K$. That is,

$$\mathcal{O}_L = \mathcal{O}_K \oplus \mathcal{O}_K\zeta_n.$$

Suppose that F/E is an extension of number fields, and let $\mathfrak{p} \leq \mathcal{O}_F$ and $\mathfrak{q} \leq \mathcal{O}_E$ be prime ideals. We say that $\mathfrak{p}$ *lies over* $\mathfrak{q}$ if $\mathfrak{p} \cap \mathcal{O}_E = \mathfrak{q}$. Given a rational prime p, one can use the Dedekind-Kummer theorem to determine the prime ideals of $\mathcal{O}_L$ lying over p. In particular, if $p \equiv 1 \bmod n$, we find that p *splits completely* in $\mathcal{O}_L$. Although this is a rather standard fact, it lies at the very core of SOLILOQUY; we therefore provide a brief explanation. Let $g \in \mathbb{F}_p$ be a primitive root. Since $n \mid p-1$, we can exhibit the element

$$w := g^{\frac{p-1}{n}} \in \mathbb{F}_p,$$

which has order n. Recall that the minimal polynomial of ζ_n over $\mathbb{Q}$ is the n-th cyclotomic polynomial $\Phi_n \in \mathbb{Z}[X]$. Let $\overline{\Phi}_n$ denote the reduction of Φ_n modulo p. Observe that the elements

$$\{w^r : r \in (\mathbb{Z}/n\mathbb{Z})^\times\}$$

all have order n in $\mathbb{F}_p$. Therefore, they are all roots of $\overline{\Phi}_n$. Since $\overline{\Phi}_n$ has degree $\phi(n)$, these elements are precisely the roots of $\overline{\Phi}_n$ in $\overline{\mathbb{F}}_p$; we deduce that $\overline{\Phi}_n$ splits into $\phi(n)$ distinct linear factors over $\mathbb{F}_p$. Since

$$p \nmid [\mathcal{O}_L : \mathbb{Z}[\zeta_n]] = 1,$$

we may apply the Dedekind-Kummer theorem; we find that p splits completely in $\mathcal{O}_L$, and the prime ideals of $\mathcal{O}_L$ lying over p are exactly those of the form

$$(p, \zeta_n - w^r),$$

where $r \in (\mathbb{Z}/n\mathbb{Z})^\times$.

2.4 Projective Space

Suppose that V is a vector space over a field F. We can define an equivalence relation $\sim$ on $V\backslash\{\mathbf{0}\}$ as follows:

$$x \sim y \iff \exists \lambda \in F^\times \text{ s.t. } x = \lambda y,$$

for $x, y \in V\backslash\{\mathbf{0}\}$. If $V = F^{r+1}$ for some positive integer r, then we will abbreviate

$$\mathbb{P}^r(F) := \mathbb{P}(V).$$

2.5 Quaternions

Let F be a totally real number field, and suppose that $a, b \in F$. The *quaternion algebra*

$$\mathcal{A} = \left(\frac{a, b}{F}\right),$$

is the algebra over F generated by elements i, j, k satisfying the following relations: $i^2 = a$, $j^2 = b$, and $ij = -ji = k$. Any element $q \in \mathcal{A}$ has a unique expression of the form

$$q = q_1 + q_2 i + q_3 j + q_4 k,$$

where $q_1, q_2, q_3, q_4 \in F$. We define the *conjugate* of q as

$$\overline{q} = q_1 - q_2 i - q_3 j - q_4 k.$$

We may further define the *reduced norm* of q as $\mathrm{nrd}(q) = q\overline{q}$, and the *reduced trace* of q as $\mathrm{trd}(q) = q + \overline{q}$.

We now make a sequence of definitions that allow us to describe the theory of ideals in the quaternion setting. First, we can capture the discrete nature of the ring of integers with the notion of an R-lattice:

Definition 5 (R-lattice). *Let R be a domain with field of fractions F. An R-lattice in an F-vector space V is a finitely generated R-submodule of V that spans V over F.*

Of course, the ring of integers has more structure than an arbitrary lattice. We next define a *maximal order*, which generalises the notion of the ring of integers to the quaternion setting.

Definition 6 (Maximal Order). *Let R be a domain with field of fractions F, and suppose that $\mathcal{A}$ is a quaternion algebra over F. An R-order $\mathcal{O} \subset \mathcal{A}$ is an R-lattice that is also a subring of $\mathcal{A}$. We say that an order is a maximal order if it is not properly contained in any other order of $\mathcal{A}$.*

We note that, in contrast to the number field case, a quaternion algebra may contain many different maximal orders.

Definition 7. *Suppose that $\mathcal{A}$ is a quaternion algebra, and let $\mathcal{L} \subset \mathcal{A}$ be a $\mathbb{Z}$-lattice. The right order of $\mathcal{L}$ is defined as*

$$\mathcal{O}_R(\mathcal{L}) := \{\alpha \in \mathcal{A} : \mathcal{L}\alpha \subset \mathcal{L}\}.$$

Note that $\mathcal{O}_R(\mathcal{L})$ is the largest order $\mathcal{O}$ for which $\mathcal{L}$ is a right $\mathcal{O}$-module.

Next, we define the parallel notions of integral and fractional ideals in the quaternion setting:

Definition 8 (Integral/Fractional Right Ideal). *Suppose that $\mathcal{A}$ is a quaternion algebra, and let $\mathcal{O} \subset \mathcal{A}$ be an order. A right fractional $\mathcal{O}$-ideal is a $\mathbb{Z}$-lattice $\mathcal{I} \subset \mathcal{A}$ such that $\mathcal{O} \subset \mathcal{O}_R(\mathcal{I})$. We say that $\mathcal{I}$ is principal if $\mathcal{I} = \alpha\mathcal{O}$ for some $\alpha \in \mathcal{A}$. If $\mathcal{I} \subset \mathcal{O}$, then we say that $\mathcal{I}$ is an integral right $\mathcal{O}$-ideal. When clear, we will simply refer to an integral right $\mathcal{O}$-ideal as a right $\mathcal{O}$-ideal.*

Finally, we introduce the corresponding notion of the norm of an ideal.

Definition 9. *Suppose that $\mathcal{A}$ is a quaternion algebra, and let $\mathcal{O} \subset \mathcal{A}$ be an order. Let $\mathcal{I} \subset \mathcal{A}$ be a right fractional $\mathcal{O}$-ideal. The reduced norm* $\mathrm{nrd}(\mathcal{I})$ *of $\mathcal{I}$ is the R-submodule of F generated by* $\{\mathrm{nrd}(\alpha) : \alpha \in \mathcal{I}\}$.

2.6 Mathematical Conventions

We conclude the preliminary mathematics by stating the remaining conventions used throughout the paper. These conventions are by no means standard, but they will aid in the presentation of this work. From now on, we will always take $n = 2^l$ for some integer $l \geq 2$, and we will abbreviate $N := \phi(n) = 2^{l-1}$. We define the fields $L = \mathbb{Q}(\zeta_n)$, $K = \mathbb{Q}(\zeta_n + \zeta_n^{-1})$, and take $\mathcal{A}$ to be the quaternion algebra given by

$$\mathcal{A} = \left(\frac{-1, -1}{K}\right).$$

We will often be interested in the particular order of $\mathcal{A}$ given by

$$\Lambda := \mathcal{O}_L \oplus j\mathcal{O}_L,$$

which we will call the *natural order*. We take the natural $\mathbb{Z}$-basis of Λ to be

$$\mathcal{B}^{quat} := \{1, \zeta_n, \ldots, \zeta_n^{N-1}, j, j\zeta_n, \ldots, j\zeta_n^{N-1}\}.$$

As such, any element $x \in \Lambda$ has a unique expression of the form

$$x = \sum_{r=0}^{N-1} x_r^{(1)} \zeta_n^r + j \sum_{r=0}^{N-1} x_r^{(2)} \zeta_n^r,$$

where $x_r^{(1)}, x_r^{(2)} \in \mathbb{Z}$. With x written as above, we define

$$||x||_\infty := \max\{|x_r^{(1)}|, |x_r^{(2)}| : 0 \leq r < N\}.$$

Since $\mathcal{B}^{quat}$ is also a $\mathbb{Q}$-basis for $\mathcal{A}$, we can define the infinity norm $||\cdot||_\infty$ on $\mathcal{A}$ in a completely analogous way. As a result of the quaternion setting, we will often use r, s to denote dummy variables, as opposed to i, j.

2.7 Cryptographic Background

We next define certain properties of public-key encryption (PKE) schemes. We begin with a security property:

Definition 10. *[16] Let $\Pi = (\mathbf{KeyGen}, \mathbf{Enc}, \mathbf{Dec})$ be a PKE scheme, and A be an adversary. We say Π is indistinguishable under chosen-plaintext attack (IND-CPA) if any probabilistic polynomial time adversary in the following experiment $\mathrm{PubK}_{\mathsf{A},\Pi}(n)$ has negligible advantage:*

1. **KeyGen** *is run to obtain keys (PK, SK).*
2. *Adversary A is given PK, and outputs a pair of equal-length messages m_0, m_1 in the message space.*
3. *A uniform bit $b \in \{0,1\}$ is chosen, and then a ciphertext $c \leftarrow \mathbf{Enc}_{PK}(m_b)$ is computed and given to A. We call c the challenge ciphertext.*
4. *A outputs a bit b'. The output of the experiment is 1 if $b' = b$, and 0 otherwise. If $b' = b$ we say that A succeeds.*

That is, $\Pr[\mathrm{PubK}_{\mathsf{A},\Pi}(n) = 1] \leq \frac{1}{2} + \mathrm{neg}(n)$.

A PKE scheme is said to be *additively homomorphic* if

$$\mathbf{Enc}_{PK}(m_0 + m_1) = \mathbf{Enc}_{PK}(m_0) + \mathbf{Enc}_{PK}(m_1).$$

3 Soliloquy

At a high level, SoliloQuat is very similar to Soliloquy, the crucial difference being that SoliloQuat relies on the principal ideal problem (PIP) for a *quaternion algebra*, as opposed to a cyclotomic number field. In this section, we give a bird's-eye view of Soliloquy, providing a list of key ingredients needed for a quaternion-based variant. This will also enable us to compare SoliloQuat to Soliloquy throughout the rest of the paper, allowing us to avoid the quantum attack exposed by Campbell-Groves-Shepherd in [4]. We have taken some artistic license in our account of Soliloquy; the original sketch of Soliloquy [4] provides very few concrete details, instead focusing on a quantum attack on the primitive. In particular, we have chosen to phrase Soliloquy as a PKE scheme (as opposed to a KEM), and have taken some inspiration from the closely-related scheme given in [26].

Key Generation A valid Soliloquy private key is an element $\alpha \in \mathcal{O}_L$ with $N_{L/\mathbb{Q}}(\alpha) = p$ for some rational prime p congruent to 1 modulo n. During key generation, we must also check that $w := 2^{(p-1)/n}$ has order n modulo p; granting this, $\alpha\mathcal{O}_L$ is a prime ideal of $\mathcal{O}_L$ lying over p, and is therefore of the form

$$(p, \zeta_n - w^r),$$

for some $r \in (\mathbb{Z}/n\mathbb{Z})^\times$. By replacing α with a suitable Galois conjugate, we may assume that $\alpha\mathcal{O}_L = (p, \zeta_n - w)$.

Encryption Recall that we take $\mathcal{B} = \{\zeta_n^r : 0 \leq r < N\}$ to be the natural integral basis for $\mathcal{O}_L$. We define the ∞-norm of an element $\varepsilon \in \mathcal{O}_L$ with respect to this basis $\mathcal{B}$. An element $\varepsilon \in \mathcal{O}_L$ is considered a *valid plaintext* if $||\varepsilon||_\infty$ is sufficiently small. In this case, we encrypt ε in two stages. First, we form a *modified plaintext* $\varepsilon' := \varepsilon + \delta$ by adding some random noise $\delta \in 2\mathcal{O}_L$. Next, we reduce modulo $\alpha\mathcal{O}_L$; since $\alpha\mathcal{O}_L = (p, \zeta_n - w)$ takes a particularly nice form, we can explicitly find an integer $0 \leq z < p$ such that

$$\varepsilon' \equiv z \bmod \alpha\mathcal{O}_L.$$

We say that z is a *valid encryption* of ε.

Decryption Suppose that we are given a valid encryption $0 \leq z < p$ of some $\varepsilon \in \mathcal{O}_L$. By construction, there exists $\delta \in 2\mathcal{O}_L$ such that

$$z \equiv \varepsilon + \delta := \varepsilon' \bmod \alpha\mathcal{O}_L.$$

In other words, there exists $x \in \mathcal{O}_L$ such that $\varepsilon' = z + \alpha x$. Rearranging, we find

$$x = \alpha^{-1}\varepsilon' - \alpha^{-1}z.$$

If $\lfloor \alpha^{-1}\varepsilon' \rceil = 0$, then we may recover x as

$$x = \lfloor -\alpha^{-1}z \rceil.$$

Therefore, $\varepsilon' = z - \alpha\lfloor \alpha^{-1}z \rceil$. Finally, since $\varepsilon \equiv \varepsilon' \bmod 2\mathcal{O}_L$, we can recover ε by reducing each $\mathcal{B}$-coefficient of ε modulo 2.

Remark 1. Importantly, we only decrypt correctly assuming that $\lfloor \alpha^{-1}\varepsilon' \rceil = 0$. One can obtain a more practical (sufficient) condition for correct decryption of the form $||\varepsilon'||_\infty < R_{Dec}$, where R_{Dec} is defined in terms of α. This was first observed by Smart and Vercauteren in [26].

Soliloquy is possible due to two key facts:

1. There is an *explicit* description of the maximal right ideals of $\mathcal{O}_L$ lying over p.
2. One can rephrase the rounding condition $\lfloor \alpha^{-1}\varepsilon' \rceil = 0$ in terms of the *radius of correct decryption* R_{Dec}.

Both of these facts can be lifted to the quaternion setting; this is the main focus of Sections 4, 5, and 6.

4 Maximal Right Ideals in the Natural Order

Throughout this section, we take p to be a fixed rational prime congruent to 1 modulo n. Since $n \mid p-1$, there exists an integer $w \in \mathbb{Z}$ with order n modulo p. For notational convenience, we define $\nu := w^{n/4}$. We will denote $\mathfrak{p}_1 := (p, \zeta_n - w)$ and $\mathfrak{p}_2 := (p, \zeta_n - w^{-1})$, which are both prime ideals of $\mathcal{O}_L$ lying over p (see the end of Section 2.3). In order to obtain a quaternion-analogue of SOLILOQUY, we need an explicit description of the maximal right ideals of Λ lying over $\mathfrak{p} := \mathfrak{p}_1 \cap \mathcal{O}_K$. This is achieved by the following proposition:

Proposition 1. *Define*

$$\mathfrak{m}_r := \begin{cases} \mathfrak{p}\Lambda + (1 - rj)(\nu + i)\Lambda & 1 \leq r \leq p, \\ \mathfrak{p}\Lambda + (\nu - i)(1 + j)\Lambda & r = p + 1. \end{cases}$$

The right ideals $\mathfrak{m}_r$ are precisely the maximal right ideals of Λ lying over $\mathfrak{p}$.

Remark 2. Taking $r = 0$, we obtain

$$\mathfrak{m}_0 = \mathfrak{p}\Lambda + (\nu + i)\Lambda = (\mathfrak{p}, \nu + i)\Lambda.$$

Moreover, working modulo $\mathfrak{p}_2$, we find that

$$\begin{aligned} \nu + i &= w^{n/4} + \zeta_n^{n/4} \\ &\equiv w^{n/4} + w^{-n/4} \\ &\equiv \nu + \nu^{-1} \equiv 0. \end{aligned}$$

In other words, $\nu + i \in \mathfrak{p}_2$. We deduce that

$$\mathfrak{m}_0 = (\mathfrak{p}, \nu + i)\Lambda \subset \mathfrak{p}_2\Lambda.$$

On the other hand, $\mathfrak{p}_2\Lambda$ is a proper right ideal of Λ, and $\mathfrak{m}_0$ is maximal. We conclude that $\mathfrak{m}_0 = \mathfrak{p}_2\Lambda$. We can generalise this observation; now consider $\mathfrak{m}_r$ for any $1 \leq r \leq p$, avoiding the case $r = p + 1$ for simplicity. We have

$$(1 - rj)(\nu + i) = \nu + (\underbrace{i - r\nu j + rk}_{\theta}).$$

Moreover, $\theta^2 = -1 - r^2(\nu^2 + 1)$. It follows that

$$F := K(\theta) \cong K\left(\sqrt{-1 - r^2(\nu^2 + 1)}\right).$$

We claim that $\mathfrak{m}_r = (\mathfrak{q} \cap \mathcal{O}_K[\theta])\Lambda$ for some prime ideal $\mathfrak{q} \leq \mathcal{O}_F$; a proof is given in Appendix A.1. We can derive some intuition for Proposition 1 from this fact. In the quaternion setting, we fail to obtain a *direct* analogue of the Dedekind-Kummer theorem. Indeed, since any element $x \in \mathcal{A}$ satisfies the quadratic

$$x^2 - \mathrm{trd}(x)x + \mathrm{nrd}(x) = 0,$$

we see that $\mathcal{A}$ is a union of infinitely many quadratic extensions of K. Moreover, if F is one of these quadratic extensions, then applying the Dedekind-Kummer theorem to $\mathcal{O}_K \subset \mathcal{O}_F$, we find a prime ideal $\mathfrak{q} \leq \mathcal{O}_F$, and $\mathfrak{q}$ corresponds to the right ideal

$$(\mathfrak{q} \cap \mathcal{O}_K[\theta])\Lambda \leq \Lambda.$$

As a result, there are *many* maximal right ideals of Λ lying over $\mathfrak{p}$. Proposition 1 demonstrates that we can construct all maximal right ideals of Λ in this way. We will revisit this remark when evaluating the security of SoliloQuat in Section 7.

Proof of Proposition 1. For clarity, we break the proof into the following steps:

1. We define a suborder $\Delta \subset \Lambda$ and a map $\chi : \Delta \to M_2(\mathbb{F}_p)$.
2. We show that χ is a surjective ring homomorphism with kernel $\mathfrak{p}\Delta$.
3. We extend χ to a ring homomorphism $\Psi : \Lambda \to M_2(\mathbb{F}_p)$ with kernel $\mathfrak{p}\Lambda$. This factors to give an explicit isomorphism $\overline{\Psi} : \Lambda/\mathfrak{p}\Lambda \to M_2(\mathbb{F}_p)$.
4. We classify the maximal right ideals $\mathcal{I}_\bullet$ of $M_2(\mathbb{F}_p)$.
5. Finally, we check that $\overline{\Psi}(\mathfrak{m}_r \bmod \mathfrak{p}\Lambda) = \mathcal{I}_\bullet$. Since $\overline{\Psi}$ gives a one-to-one correspondence between maximal right ideals of Λ lying over $\mathfrak{p}$, and maximal right ideals of $M_2(\mathbb{F}_p)$, we deduce the result.

Step 1: We define the suborder

$$\Delta := \mathcal{O}_K \oplus i\mathcal{O}_K \oplus j\mathcal{O}_K \oplus k\mathcal{O}_K \subset \Lambda.$$

By definition, any element $x \in \Delta$ can be written uniquely in the form

$$x = a + bi + cj + dk,$$

where $a, b, c, d \in \mathcal{O}_K$. Since p splits completely in $\mathcal{O}_L$, we know that $\mathcal{O}_K/\mathfrak{p} \cong \mathbb{F}_p$. We can therefore define $\chi(x)$ as the matrix

$$\begin{pmatrix} a + b\nu & c + d\nu \\ -c + d\nu & a - b\nu \end{pmatrix} \in M_2(\mathbb{F}_p),$$

where each entry is taken modulo $\mathfrak{p}$.

Step 2: Consider the following linear system over $\mathcal{O}_K/\mathfrak{p} \cong \mathbb{F}_p$:

$$\underbrace{\begin{pmatrix} 1 & \nu & 0 & 0 \\ 0 & 0 & 1 & \nu \\ 0 & 0 & -1 & \nu \\ 1 & -\nu & 0 & 0 \end{pmatrix}}_{M} \begin{pmatrix} a \\ b \\ c \\ d \end{pmatrix} = \begin{pmatrix} a + b\nu \\ c + d\nu \\ -c + d\nu \\ a - b\nu \end{pmatrix}. \tag{4}$$

Since M has full rank, we deduce that χ is surjective. Next, we verify that χ is a ring homomorphism; suppose that we are given another element

$$x' = a' + b'i + c'j + d'k \in \Delta,$$

where $a', b', c', d' \in \mathcal{O}_K$. We find that

$$\begin{aligned} xx' &= (a + bi + cj + dk)(a' + b'i + c'j + d'k) \\ &= (aa' - bb' - cc' - dd') + (ab' + a'b + cd' - c'd)i \\ &\qquad + (ac' + a'c + b'd - bd')j + (ad' + a'd + bc' - b'c)k. \end{aligned}$$

Therefore, $\chi(xx')$ is given by the matrix

$$\begin{pmatrix} (aa'-bb'-cc'-dd') & (ac'+a'c+b'd-bd') \\ +(ab'+a'b+cd'-c'd)\nu & +(ad'+a'd+bc'-b'c)\nu \\ -(ac'+a'c+b'd-bd') & (aa'-bb'-cc'-dd') \\ +(ad'+a'd+bc'-b'c)\nu & -(ab'+a'b+cd'-c'd)\nu \end{pmatrix},$$

where every entry is taken modulo $\mathfrak{p}$. On the other hand, $\chi(x)\chi(x')$ is given by

$$\begin{pmatrix} a+b\nu & c+d\nu \\ -c+d\nu & a-b\nu \end{pmatrix}\begin{pmatrix} a'+b'\nu & c'+d'\nu \\ -c'+d'\nu & a'-b'\nu \end{pmatrix},$$

where again, each entry is taken modulo $\mathfrak{p}$. Multiplying these matrices directly, we find that $\chi(xx') = \chi(x)\chi(x')$. Moreover,

$$\begin{aligned}\chi(x)+\chi(x') &= \begin{pmatrix} a+b\nu & c+d\nu \\ -c+d\nu & a-b\nu \end{pmatrix} + \begin{pmatrix} a'+b'\nu & c'+d'\nu \\ -c'+d'\nu & a'-b'\nu \end{pmatrix} \\ &= \begin{pmatrix} (a+a')+(b+b')\nu & (c+c')+(d+d')\nu \\ -(c+c')+(d+d')\nu & (a+a')-(b+b')\nu \end{pmatrix} = \chi(x+x').\end{aligned}$$

Therefore, χ is a surjective ring homomorphism. Finally, we check that χ has kernel $\mathfrak{p}\Delta$. By linear system (4), we see that $\chi(x) = 0$ precisely when

$$a \equiv b \equiv c \equiv d \bmod \mathfrak{p}.$$

In other words, $\chi(x) = 0$ exactly when $x \in \mathfrak{p}\Delta$, as required.

Step 3: We now extend χ to Λ. Take an arbitrary element $y \in \Lambda$. Since $\{1, \zeta_n\}$ is a relative integral basis for $\mathcal{O}_L/\mathcal{O}_K$, we may write

$$y = e + f\zeta_n + gj + hj\zeta_n,$$

where $e, f, g, h \in \mathcal{O}_K$. Observe that

$$2(e+f\zeta_n) = 2e + f(\zeta_n+\zeta_n^{-1}) + f\left(\frac{\zeta_n-\zeta_n^{-1}}{i}\right)i \in \mathcal{O}_K[i]. \tag{5}$$

We deduce that $2y \in \Delta$. This allows us to define $\Psi : \Lambda \to M_2(\mathbb{F}_p)$ as

$$\Psi(y) = \begin{pmatrix} 2^{-1} & 0 \\ 0 & 2^{-1} \end{pmatrix}\chi(2y).$$

Surjectivity of Ψ is immediate since Ψ and χ agree on Δ. Furthermore, we can use the fact that χ is a ring homomorphism to show that Ψ is also a ring homomorphism. Taking $y, y' \in \Lambda$, we compute

$$\begin{aligned}\Psi(yy') &= \begin{pmatrix} 2^{-1} & 0 \\ 0 & 2^{-1} \end{pmatrix}\chi(2yy') \\ &= \begin{pmatrix} 2^{-1} & 0 \\ 0 & 2^{-1} \end{pmatrix}^2\chi(2y)\chi(2y') \\ &= \begin{pmatrix} 2^{-1} & 0 \\ 0 & 2^{-1} \end{pmatrix}\chi(2y)\begin{pmatrix} 2^{-1} & 0 \\ 0 & 2^{-1} \end{pmatrix}\chi(2y') = \Psi(y)\Psi(y').\end{aligned}$$

Similarly, we find that $\Psi(y) + \Psi(y') = \Psi(y + y')$. It remains to show that Ψ has kernel $\mathfrak{p}\Lambda$. Since χ has kernel $\mathfrak{p}\Delta$, we find that

$$\ker(\Psi) = \{y \in \Lambda : 2y \in \mathfrak{p}\Delta\}.$$

In order to conclude that $\ker(\Psi) = \mathfrak{p}\Lambda$, we must show that, for $y \in \Lambda$,

$$2y \in \mathfrak{p}\Delta \iff y \in \mathfrak{p}\Lambda.$$

Clearly, if $y \in \mathfrak{p}\Lambda$, then $2y \in \mathfrak{p}(2\Lambda) \subset \mathfrak{p}\Delta$. Conversely, suppose that $y \in \Lambda$ and $2y \in \mathfrak{p}\Delta$. By equation (5), we find that

$$2e + f(\zeta_n + \zeta_n^{-1}),\ f\left(\frac{\zeta_n - \zeta_n^{-1}}{i}\right) \in \mathfrak{p}.$$

Since $\zeta_n - \zeta_n^{-1} \notin \mathfrak{p}$ and $\mathfrak{p}$ is a prime ideal, we deduce that $f \in \mathfrak{p}$. Therefore, $2e \in \mathfrak{p}$, and since $2 \notin \mathfrak{p}$, we see that $e \in \mathfrak{p}$ also. An identical argument shows that $g, h \in \mathfrak{p}$; it follows that $y \in \mathfrak{p}\Lambda$. We deduce that $\ker(\Psi) = \mathfrak{p}\Lambda$.

Step 4: We now describe the maximal right ideals of $M_2(\mathbb{F}_p)$. Suppose that $\mathcal{I}$ is a proper right ideal of $M_2(\mathbb{F}_p)$. Since $\mathcal{I}$ is proper, $\mathcal{I}$ cannot contain an invertible matrix. In other words, every non-zero matrix in $\mathcal{I}$ has rank one. Let $\mathbf{M} \in \mathcal{I}\backslash\{\mathbf{0}\}$ have rows $\mathbf{r}_1, \mathbf{r}_2$; since $\mathbf{M}$ has rank one, $\mathbf{r}_1$ and $\mathbf{r}_2$ are linearly dependent. Therefore, it will be useful to define

$$\mathbf{r}_2/\mathbf{r}_1 = \begin{cases} \infty & \mathbf{r}_1 = 0, \\ (\mathbf{r}_2)_1/(\mathbf{r}_1)_1 & (\mathbf{r}_1)_1 \neq 0, \\ (\mathbf{r}_2)_2/(\mathbf{r}_1)_2 & (\mathbf{r}_1)_2 \neq 0. \end{cases}$$

In particular, let $\mathbf{r}_2/\mathbf{r}_1 = \lambda \in \mathbb{P}^1(\mathbb{F}_p)$. Since the natural right-action of $M_2(\mathbb{F}_p)$ on $\mathbb{F}_p^2\backslash\{\mathbf{0}\}$ is transitive, we find that

$$\mathcal{I} = \mathcal{I}_\lambda := \left\{ A = \begin{pmatrix} - & \mathbf{r}_1 & - \\ - & \mathbf{r}_2 & - \end{pmatrix} : \mathbf{r}_1, \mathbf{r}_2 \in \mathbb{F}_p^2,\ \mathbf{r}_2/\mathbf{r}_1 = \lambda \right\}.$$

We deduce that the maximal right ideals of $M_2(\mathbb{F}_p)$ are precisely those of the form $\mathcal{I}_\lambda$ where $\lambda \in \mathbb{P}^1(\mathbb{F}_p)$.

Step 5: Finally, we must check that

$$\{\Psi(\mathfrak{m}_r) : 1 \leq r \leq p+1\} = \{\mathcal{I}_\lambda : \lambda \in \mathbb{P}^1(\mathbb{F}_p)\}.$$

Indeed, we find that for $1 \leq r \leq p$,

$$\Psi(\mathfrak{m}_r) = \begin{pmatrix} 2\nu & 0 \\ 2r\nu & 0 \end{pmatrix} M_2(\mathbb{F}_p).$$

On the other hand, if $r = p + 1$, we have

$$\Psi(\mathfrak{m}_r) = \begin{pmatrix} 0 & 0 \\ -2\nu & 2\nu \end{pmatrix} M_2(\mathbb{F}_p).$$

We deduce that

$$\Psi(\mathfrak{m}_r) = \begin{cases} \mathcal{I}_r & 1 \leq r \leq p, \\ \mathcal{I}_\infty & r = p + 1. \end{cases}$$

□

5 SoliloQuat

We are now ready to detail SoliloQuat, which consists of three main components: key generation, encryption, and decryption. While the parameters chosen during key generation may seem rather arbitrary, these choices will become transparent when we discuss correctness.

Key Generation A candidate SoliloQuat private key is an element of the form

$$\alpha = \sum_{r=0}^{N-1} (a_r^{(1)} + j a_r^{(2)}) \zeta_n^r \in \Lambda,$$

where $a_r^{(1)}, a_r^{(2)} \in \mathbb{Z}$. We will generate such a candidate by sampling each $\mathcal{B}^{quat}$-coefficient $a_r^{(1)}, a_r^{(2)}$ from a discrete Gaussian of mean 0 and width $\sigma_N > O(N^{5/4})$. Before we accept α as a *valid SoliloQuat private key*, we must check that α has certain properties, the first two of which are as follows:

C1. $|N_{K/\mathbb{Q}}(\mathrm{nrd}(\alpha))| = p$ for some rational prime p with $p \equiv 1 \bmod n$.
C2. $w := 2^{(p-1)/n}$ has order n modulo p.

These two conditions C1-2 can be seen as a direct analogue of those required for a valid Soliloquy private key, and will give us a convenient representation of the right ideal $\alpha\Lambda \leq \Lambda$. Indeed, if α satisfies property C1, then $\alpha\Lambda$ is a maximal right ideal lying over some prime ideal $\mathfrak{p} \leq \mathcal{O}_K$ with $\mathfrak{p} \cap \mathbb{Z} = (p)$. Assuming property C2, we have

$$\mathfrak{p} = (p, \zeta_n - w^r) \cap \mathcal{O}_K,$$

for some $r \in (\mathbb{Z}/n\mathbb{Z})^\times$. In order to assume that $r = 1$, we define a natural action of $\mathrm{Gal}(L/\mathbb{Q})$ on Λ. Suppose that we are given an element

$$x = a + jb \in \Lambda,$$

where $a, b \in \mathcal{O}_L$. Given $\sigma \in \mathrm{Gal}(L/\mathbb{Q})$, we simply define

$$\sigma(x) = \sigma(a) + j\sigma(b).$$

Crucially, this action is compatible with reduced norms:

$$\begin{aligned}\mathrm{nrd}(\sigma(x)) &= \sigma(a)\overline{\sigma(a)} + \sigma(b)\overline{\sigma(b)} \\ &= \sigma(a\overline{a}) + \sigma(b\bar{b}) \\ &= \sigma(a\overline{a} + b\bar{b}) = \sigma(\mathrm{nrd}(x)).\end{aligned}$$

Therefore, taking a suitable Galois conjugate of α, we may assume that $\alpha\Lambda$ lies over $\mathfrak{p} = (p, \zeta_n - w) \cap \mathcal{O}_K$. By Proposition 1, we have $\alpha\Lambda = \mathfrak{m}_r$ for some $1 \leq r \leq p+1$.

We are now in a very similar setting to that of SOLILOQUY; we must take some care to avoid the quantum attack presented in [4]. Following Remark 2, consider the number field

$$F = K(\underbrace{i - r\nu j + rk}_{\theta}) \cong K\left(\sqrt{-1 - r^2(\nu^2+1)}\right).$$

In Appendix A.1, we show that $\mathfrak{m}_r$ is of the form $(\mathfrak{q} \cap \mathcal{O}_K[\theta])\Lambda$, where $\mathfrak{q} \leq \mathcal{O}_F$ is a prime ideal lying over $\mathfrak{p}$. If $\mathfrak{q} \cap \mathcal{O}_K[\theta]$ is principal, one can use the Lenstra-Silverberg algorithm [20] to find a principal generator of $\mathfrak{q} \cap \mathcal{O}_K[\theta]$, which is automatically a principal generator of $\mathfrak{m}_r$. This could lead to a direct key recovery attack on SOLILOQUAT. In order to avoid this, we impose a *third* condition on a valid SOLILOQUAT private key:

C3. Let $\mathfrak{q}_1, \mathfrak{q}_2$ be the prime ideals of $\mathcal{O}_F$ lying over $\mathfrak{p}$. Ensure that $\mathfrak{q}_1 \cap \mathcal{O}_K[\theta]$ and $\mathfrak{q}_2 \cap \mathcal{O}_K[\theta]$ are not principal ideals of $\mathcal{O}_K[\theta]$.

We can verify condition C3 in deterministic polynomial time using the Lenstra-Silverberg algorithm; further details are given in Appendix A.2. If $\alpha \in \Lambda$ satisfies conditions C1-3, then α is a *valid SoliloQuat private key*.

Remark 3. We will later analyse the *radius of correct decryption*, which we will define as a random variable R_{Dec} depending on the (random) private key α. One could add an extra step in the **KeyGen** algorithm to check that the radius of correct decryption is above a certain threshold. We have omitted this step as we have not yet defined R_{Dec}.

Remark 4. When checking that $\alpha \in \Lambda$ is a valid SOLILOQUAT private key, Condition C2 holds with high probability. For example, $2^{(p-1)/n}$ has order n modulo p whenever 2 is a primitive root. Under GRH, this occurs with probability ≈ 0.37396 [27,15]. Therefore, the 'main' condition is that the absolute norm of α is prime. We conjecture that condition C3 also holds with high probability; in order to avoid wading through class group heuristics, we simply check this condition during key generation.

Algorithm 1: KeyGen(σ_N)

Input: – $\sigma_N \in \mathbb{R}^+$ with $\sigma_N > O(N^{5/4})$.
– Access to an oracle **IsPrincipal**: given a CM-order A and an ideal $\mathcal{I} \leq A$, **IsPrincipal** returns 'true' if $\mathcal{I}$ is principal, and 'false' otherwise.
– Access to an oracle **IsPrime**: given an integer $n \in \mathbb{Z}$, **IsPrime** returns 'true' if n is prime, and 'false' otherwise.

Output: A valid SOLILOQUAT private key SK with corresponding public key PK.

done $\leftarrow$ false
while not done **do**
  $a_r^{(1)}, a_r^{(2)} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma_N}$ for $0 \leq r < N$
  $\alpha \leftarrow \sum_{r=0}^{N-1} a_r^{(1)} \zeta_n^r + j \sum_{r=0}^{N-1} a_r^{(2)} \zeta_n^r$
  $p \leftarrow N_{K/\mathbb{Q}}(\mathrm{nrd}(\alpha))$
  PrimeNorm$\leftarrow$**IsPrime**(p)
  if PrimeNorm=false **or** $p \not\equiv 1 \bmod n$ **or** $2^{\frac{p-1}{2}} \equiv 1 \bmod p$ **then**
    continue
  end if
  Let $w \in \{1, 2, \ldots, p\}$ have $w \equiv 2^{(p-1)/n} \bmod p$
  Let $\nu \in \{1, 2, \ldots, p\}$ have $\nu \equiv w^{n/4} \bmod p$
  $\mathfrak{p} \leftarrow \mathrm{nrd}(\alpha)\mathcal{O}_K$
  $\mathbf{M} \leftarrow \Psi(\alpha)$
  $\mathbf{r}_1 \leftarrow \mathbf{M}_{1,\bullet}$
  $\mathbf{r}_2 \leftarrow \mathbf{M}_{2,\bullet}$
  if $\mathbf{r}_1 \neq 0$ **then**
    Let $s \in \{1, 2, \ldots, p\}$ have $s \equiv \mathbf{r_2}/\mathbf{r}_1 \bmod p$
    $\theta \leftarrow (1 - sj)(\nu + i)$
  else
    $s \leftarrow p + 1$
    $\theta \leftarrow (\nu - i)(1 + j)$
  end if
  $A \leftarrow \mathcal{O}_K[\theta]$
  $F \leftarrow K(\theta)$
  if not IsPrincipal(A, $(\mathfrak{p}\mathcal{O}_F + \theta\mathcal{O}_F) \cap A$) **then**
    done $\leftarrow$ true
  end if
end while
$SK \leftarrow \alpha$
$PK \leftarrow (\mathfrak{p}, \theta)$
return (SK, PK)

Encryption Suppose that α is a valid SoliloQuat private key with corresponding public key $PK = (\mathfrak{p}, \theta)$. We define a *valid plaintext* to be an element $\varepsilon \in \Lambda$ of the form

$$\varepsilon = \sum_{r=0}^{N-1} (e_r^{(1)} + j e_r^{(1)}) \zeta_n^r \in \Lambda,$$

where $e_r^{(1)}, e_r^{(2)} \in \{0, 1\}$. Mirroring Soliloquy, we will encrypt ε in two stages:

1. Form a *modified plaintext* $\varepsilon' := \varepsilon + \delta$, where δ is some random noise. Since we take $\delta \in 2\Lambda$, we will refer to δ as *even noise*.

2. Reduce ε' modulo $\mathfrak{p}\Lambda + \theta\Lambda$.

We begin with a description of stage 1. First, choose $\sigma'_N < O(\sigma_N / N^{5/4})$ and sample

$$\delta' = \sum_{r=0}^{N-1} d_r^{(1)} \zeta_n^r + j \sum_{r=0}^{N-1} d_r^{(2)} \zeta_n^r,$$

where $d_r^{(1)}, d_r^{(2)} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma'_N/2}$. We then define $\delta := 2\delta'$ and form the element $\varepsilon' = \varepsilon + \delta$. Crucially, we can recover ε from *any* modified plaintext ε' using the fact that

$$\varepsilon' = \varepsilon + 2\delta' \equiv \varepsilon \bmod 2\Lambda.$$

In other words, we can recover ε by reducing each $\mathcal{B}^{quat}$-coefficient of ε' modulo 2. Note that the correctness of SoliloQuat relies on σ'_N being sufficiently small; this will be made into a rigorous statement in Section 6. We now give a concrete description of the second stage of encryption, that is, reducing modulo $\mathfrak{p}\Lambda + \theta\Lambda$:

1. Compute $1 \leq s \leq p + 1$ such that $\mathfrak{m}_s = \mathfrak{p}\Lambda + \theta\Lambda$.

2. Define the following matrix over $\mathcal{O}_K / \mathfrak{p} \cong \mathbb{F}_p$:

$$A_s = \begin{cases} \begin{pmatrix} 1 & -s^{-1} \\ 0 & 0 \end{pmatrix}, & \text{if } 1 \leq s < p, \\ \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}, & \text{if } s = p, \\ \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, & \text{if } s = p + 1. \end{cases}$$

   Following Proposition 1, one can easily see that $\Psi(\varepsilon')$ and $A_s\Psi(\varepsilon')$ lie in the same class modulo $\Psi(\mathfrak{m}_s)$. For example, if $1 \leq s < p$ we have

$$A_s\Psi(\varepsilon') = I_2\Psi(\varepsilon') - \begin{pmatrix} 0 & s^{-1} \\ 0 & 1 \end{pmatrix} \Psi(\varepsilon') \equiv \Psi(\varepsilon') \bmod \Psi(\mathfrak{m}_s).$$

Write

$$A_s\Psi(\varepsilon') = \begin{pmatrix} a & b \\ c & d \end{pmatrix},$$

where $a, b, c, d \in \mathcal{O}_K/\mathfrak{p}$.

3. As $\mathcal{O}_K/\mathfrak{p} \cong \mathbb{F}_p$, there are integers $0 \leq z_a, z_b, z_c, z_d < p$ such that

$$z_a \equiv a \bmod \mathfrak{p}, \quad z_b \equiv b \bmod \mathfrak{p},$$
$$z_c \equiv c \bmod \mathfrak{p}, \quad z_d \equiv d \bmod \mathfrak{p}.$$

Again, since $\mathcal{O}_K/\mathfrak{p} \cong \mathbb{F}_p$, we can identify z_a, z_b, z_c, z_d with elements of $\mathbb{F}_p$, and since $\mathfrak{p} = (\zeta_n + \zeta_n^{-1}, w - w^{-1})$, we can efficiently compute z_a, z_b, z_c, z_d. For an integer $x \in \{1, 2, \ldots, p-1\}$, let $0 \leq x_{inv} < p$ be the unique integer with $xx_{inv} \equiv 1 \bmod p$. We note that

$$\begin{aligned}\varepsilon' \equiv z := {} & 2_{inv}(z_a + z_d) + (2\nu)_{inv}(z_a - z_d)i \\ & + 2_{inv}(z_b - z_c)j + (2\nu)_{inv}(z_b + z_c)k \pmod{\mathfrak{m}_s}.\end{aligned}$$

We will choose this expression as a reduction of ε' modulo $\mathfrak{m}_s = \mathfrak{p}\Lambda + \theta\Lambda$. Although we may associate z with a quaternion in $(-1, -1 \,|\, \mathbb{F}_p)$,[6] we will find it more convenient to view z as an element of Λ.

A concrete description of SoliloQuat encryption is given in Algorithm 2 below.

Decryption Let α be a valid SoliloQuat private key. Suppose that we are given a valid SoliloQuat encryption $z \in \Lambda$ of $\varepsilon \in \Lambda$. By construction, there exists $\delta' \in \Lambda$ such that

$$\varepsilon' := \varepsilon + 2\delta' \equiv z \bmod \alpha\Lambda.$$

Therefore, there exists $x \in \Lambda$ such that $\varepsilon' - z = \alpha x$. Rearranging, we find that

$$\alpha^{-1}\varepsilon' - \alpha^{-1}z = x.$$

Let $\lfloor \cdot \rceil$ denote coordinate-wise rounding with respect to $\mathcal{B}^{quat}$. Assuming that $\lfloor \alpha^{-1}\varepsilon' \rceil = 0$, we find that $x = \lfloor -\alpha^{-1}z \rceil$. Therefore, we may recover ε' as

$$\varepsilon' = z - \alpha\lfloor \alpha^{-1}z \rceil.$$

As before, we can find ε by reducing each $\mathcal{B}^{quat}$-coefficient of ε' modulo 2.

A concrete description of SoliloQuat decryption is given in Algorithm 3 below.

[6] Since $A_s\Psi(\varepsilon')$ has only two non-zero entries, we can even associate z with an element of $\mathbb{F}_p^2$, which is consistent with the fact that $\Lambda/\alpha\Lambda \cong \mathbb{F}_p^2$ as additive groups.

Algorithm 2: Enc(ε, σ'_N, $PK = (\mathfrak{p}, \theta)$)

Input: – $\sigma'_N \in \mathbb{R}^+$.
– $\varepsilon = \sum_{r=0}^{N-1} e_r^{(1)} \zeta_n^r + j \sum_{r=0}^{N-1} e_r^{(2)} \zeta_n^r$ where $e_r^{(1)}, e_r^{(2)} \in \{0, 1\}$.
– Let p denote the rational prime lying under $\mathfrak{p}$. For an integer $1 \leq x < p$, let $1 \leq x_{inv} < p$ be such that $x x_{inv} \equiv 1 \bmod p$.

Output: A valid SoliloQuat encryption z of ε.

$d_r^{(1)}, d_r^{(2)} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma'_N/2}$ for $0 \leq r < N$
$\delta \leftarrow 2\left(\sum_{r=0}^{N-1} d_r^{(1)} \zeta_n^r + j \sum_{r=0}^{N-1} d_r^{(2)} \zeta_n^r\right)$
$\varepsilon' \leftarrow \varepsilon + \delta$
$\mathbf{M} \leftarrow \Psi(\theta)$
$\mathbf{r}_1 \leftarrow \mathbf{M}_{1,\bullet}$
$\mathbf{r}_2 \leftarrow \mathbf{M}_{2,\bullet}$
if $\mathbf{r}_1 \neq 0$ **then**
  Let $s \in \{1, 2, \ldots, p\}$ have $s \equiv \mathbf{r_2}/\mathbf{r}_1 \bmod p$
  if $s = p$ **then**
    $A_s \leftarrow \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$
  else
    $A_s \leftarrow \begin{pmatrix} 1 & -s_{inv} \\ 0 & 0 \end{pmatrix}$
  end if
else
  $s \leftarrow p + 1$
  $A_s \leftarrow \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$
end if
$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \leftarrow A_s \Psi(\varepsilon')$, where a, b, c, d are written in terms of the power basis $\{(\zeta_n + \zeta_n^{-1})^r : 0 \leq r < N/2\}$ of $\mathcal{O}_K$.
Evaluate a, b, c, d at $\zeta_n + \zeta_n^{-1} = w + w^{-1}$ and reduce modulo p to obtain integers z_a, z_b, z_c, z_d respectively.
Let $w \in \{1, 2, \ldots, p\}$ have $w \equiv 2^{(p-1)/n} \bmod p$
Let $\nu \in \{1, 2, \ldots, p\}$ have $\nu \equiv w^{n/4} \bmod p$
$z \leftarrow 2_{inv}(z_a + z_d) + (2\nu)_{inv}(z_a - z_d)i + 2_{inv}(z_b - z_c)j + (2\nu)_{inv}(z_b + z_c)k$
return z

Algorithm 3: Dec(z, $SK = \alpha$)

Input: – A valid SoliloQuat private key $\alpha \in \Lambda$. Let $PK = (\mathfrak{p}, \theta)$ denote the corresponding public key.
– $z = \mathbf{Enc}(\varepsilon, \sigma'_N, PK)$ for some $\sigma'_N \in \mathbb{R}^+$.

Output: $\varepsilon \in \Lambda$

$\varepsilon' \leftarrow z - \alpha \lfloor \alpha^{-1} z \rceil$
$\varepsilon \leftarrow \varepsilon' \bmod 2\Lambda$
return ε

6 Correctness of Decryption

Let α be a valid SoliloQuat private key with corresponding public key $PK = (\mathfrak{p}, \theta)$, and let $\varepsilon \in \Lambda$ be a valid plaintext. Recall that SoliloQuat encryption consists of two stages:

1. Form a *modified plaintext* $\varepsilon' := \varepsilon + \delta$, where δ is some random even noise.
2. Reduce ε' modulo $\mathfrak{p}\Lambda + \theta\Lambda = \alpha\Lambda$.

Importantly, we can only decrypt correctly if $\lfloor \alpha^{-1}\varepsilon' \rceil = 0$. Unfortunately, this condition involves α and ε', both of which can be viewed as random variables, sampled during key generation and encryption respectively. As a result, the distribution of $\alpha^{-1}\varepsilon'$ is rather hard to control. Taking inspiration from [26], we will show that decryption is correct, assuming that

$$||\varepsilon'||_\infty < R_{Dec},$$

where R_{Dec} is a random variable that *depends only on* α. We will then give a probable bound on R_{Dec}, giving a transparent condition for correct decryption. Since $\mathcal{A}$ is a division algebra and $\alpha \neq 0$, we may write

$$\beta := \alpha^{-1} = \sum_{r=0}^{N-1} (b_r^{(1)} + j b_r^{(2)}) \zeta_n^r \in \mathcal{A},$$

where $b_r^{(1)}, b_r^{(2)} \in \mathbb{Q}$. Note that the condition for correct decryption $\lfloor \beta\varepsilon' \rceil = 0$ may be restated as

$$||\beta\varepsilon'||_\infty < \frac{1}{2}. \tag{6}$$

The following lemma will allow us to 'separate' β and ε':

Lemma 1. *[Appendix A.3] Suppose that $\beta, \varepsilon' \in \mathcal{A}$. Then*

$$||\beta\varepsilon'||_\infty \leq 2N||\beta||_\infty||\varepsilon'||_\infty.$$

By Lemma 1, we have correct decryption provided that

$$2N||\beta||_\infty||\varepsilon'||_\infty < \frac{1}{2}. \tag{7}$$

In order to obtain a condition for correct decryption that doesn't depend on β, we will exhibit an upper bound on $||\beta||_\infty$ that holds with overwhelming probability. For an element $q \in \mathcal{A}$, we define L_q as the map $\mathcal{A} \to \mathcal{A}$ given by $x \mapsto qx$. Throughout the rest of the paper, for any $q \in \mathcal{A}$, we take M_q be the matrix representing L_q with respect to $\mathcal{B}^{quat}$. The following lemma gives an upper bound on $||\beta||_\infty$:

Lemma 2. *[Appendix A.3] Suppose we are given an element $0 \neq \alpha \in \Lambda$ and define $\beta = \alpha^{-1}$. Let M_α have eigenvalues $\lambda_N^0, \lambda_N^1, \ldots, \lambda_N^{2N-1}$. Then*

$$||\beta||_\infty \leq \frac{1}{\min_r |\lambda_N^r|}.$$

Combining Lemma 2 and inequality (7), we see that the following condition is sufficient for correct decryption:

$$\frac{2N}{\min_r |\lambda_N^r|}||\varepsilon'||_\infty < \frac{1}{2} \iff ||\varepsilon'||_\infty < \frac{\min_r |\lambda_N^r|}{4N} := R_{Dec}. \tag{8}$$

We will refer to R_{Dec} as the *radius of correct decryption*. It is now clear that, in order to understand the radius of correct decryption, we must analyse the (random) eigenvalues λ_N^r of M_α, where α is sampled as in the **KeyGen** algorithm. To that end, we have the following result:

Corollary 1. *[Appendix A.4] Let $\sigma_N \geq 4N^{3/4}$ and sample $a_r^{(1)}, a_r^{(2)} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma_N}$ for $0 \leq r < N$. Form the element*

$$\alpha = \sum_{r=0}^{N-1} a_r^{(1)} \zeta_n^r + j \sum_{r=0}^{N-1} a_r^{(2)} \zeta_n^r.$$

Let $\{\lambda_N^r : 0 \leq r < 2N\}$ be the eigenvalues of M_α and define the random variable

$$R_{Dec} = \frac{\min_r |\lambda_N^r|}{4N}.$$

Then,

$$\mathbb{P}\left(R_{Dec} \geq \frac{\sigma_N}{4N^{3/4}}\right) \gtrsim \frac{3}{4}.$$

Recall that we take $\sigma'_N < O(\sigma_N / N^{5/4})$. In particular, let

$$\sigma'_N = \sigma_N / C N^{5/4},$$

for some large constant $C > 0$. Following Corollary 1, we must check that this choice of σ'_N ensures that

$$||\varepsilon'||_\infty < R_{Dec},$$

with overwhelming probability. We may achieve this with a simple application of Chebychev's inequality:

Proposition 2. *Suppose that $\eta_N > 0$ and let $\{x_r^{(1)}, x_r^{(2)} : 0 \leq r < N\}$ be IID with distribution $\mathcal{D}_{\mathbb{Z},\eta_N}$. Form the element*

$$x = \sum_{r=0}^{N-1} x_r^{(1)} \zeta_n^r + j \sum_{r=0}^{N-1} x_r^{(2)} \zeta_n^r.$$

For any constant $D > 1$, we have

$$\mathbb{P}\left(||x||_\infty < D\eta_N \sqrt{N}\right) \gtrsim e^{-\frac{2}{D^2}}.$$

Proof. Note that each $x_r^{(1)}, x_r^{(2)}$ has mean 0 and variance approximately η_N^2. By Chebychev's inequality, we have

$$\begin{aligned}\mathbb{P}\left(||x||_\infty < D\eta_N\sqrt{N}\right) &= \mathbb{P}\left(|x_r^{(1)}|, |x_r^{(2)}| < D\eta_N\sqrt{N} \text{ for } 0 \leq r < N\right) \\ &= \left(1 - \mathbb{P}\left(|x_1^{(1)}| \geq D\eta_N\sqrt{N}\right)\right)^n \\ &\gtrsim \left(1 - \frac{1}{D^2N}\right)^n \approx e^{-\frac{2}{D^2}}.\end{aligned}$$

□

As usual, write $\varepsilon' = \varepsilon + 2\delta'$ where we sample $\delta' \leftarrow \mathcal{D}_{\mathbb{Z},\sigma'_N/2}$ during the encryption of ε. Applying Proposition 2 with $\eta_N = \sigma'_N/2$ and some large constant $D > 1$, we find that

$$||\delta'||_\infty < \frac{1}{2}D\sigma'_N\sqrt{N} = D\sigma_N/2CN^{3/4},$$

with overwhelming probability. Taking C to be sufficiently large and applying Corollary 1, we conclude that

$$||\varepsilon'||_\infty < \frac{D\sigma_N}{CN^{3/4}} < R_{Dec},$$

with overwhelming probability. In other words, taking $\sigma'_N << \sigma_N/N^{5/4}$ gives a correct scheme.

7 Security

In order to obtain rigorous security reductions, we will recast the arithmetic of $\mathcal{A}$ in terms of lattices. After analysing the typical 'length' of a modified plaintext $\varepsilon' := \varepsilon + \delta$, we will be able to reduce the security of SoliloQuat to different forms of the *approximate bounded distance decoding problem (γ-BDD).*

Throughout this section, let $\alpha \in \Lambda$ be a valid SoliloQuat private key, with corresponding public key $PK = (\mathfrak{p}, \theta)$. We will consider the encryption of an element $\varepsilon \in \Lambda$ with a modified plaintext $\varepsilon' := \varepsilon + \delta$, where each $\mathcal{B}^{quat}$-coefficient of δ has variance around σ'^2_N. By construction, the encryption z of ε satisfies

$$z \equiv \varepsilon' \bmod \alpha\Lambda. \tag{9}$$

Taking inspiration from the *coefficient embedding* of a cyclotomic number field, we will recast the above equation in terms of lattices. Indeed, any element $x \in \Lambda$ has a unique expression of the form

$$x = \sum_{r=0}^{N-1} x_r^{(1)}\zeta_n^r + j\sum_{r=0}^{N-1} x_r^{(2)}\zeta_n^r,$$

where $x_r^{(1)}, x_r^{(2)} \in \mathbb{Z}$, so we may define an embedding $\iota : \Lambda \to \mathbb{Z}^n$ by

$$\iota(x) = \left(x_0^{(1)}\ x_1^{(1)} \cdots x_{N-1}^{(1)}\ x_0^{(2)}\ x_1^{(2)} \cdots x_{N-1}^{(2)}\right)^{\intercal}.$$

Naturally, $\iota(\alpha\Lambda)$ is a sublattice of $\mathbb{Z}^n$; we will denote this sublattice by $\mathcal{L}_\alpha$. Since M_α represents left-multiplication by α with respect to $\mathcal{B}^{quat}$, we know that $\mathcal{L}_\alpha$ has generating matrix M_α. We may rephrase equation (9) as

$$\iota(z) \equiv \iota(\varepsilon') \bmod \mathcal{L}_\alpha. \tag{10}$$

7.1 Typical Length of a Modified Plaintext

We would like to compare the lengths of $\iota(\varepsilon')$ and $\iota(\alpha)$ to $\lambda_1(\mathcal{L}_\alpha)$. To that end, we have the following result:

Proposition 3. *Let $\sigma_N > CN^{5/4}$ for some constant $C > 0$, and define*

$$\sigma'_N = \sigma_N / CN^{5/4}.$$

Suppose that we generate random elements $\alpha, \varepsilon' \in \Lambda$ according to the algorithms **KeyGen**(σ_N) *and* **Enc**$(\varepsilon, \sigma'_N, PK)$ *respectively, where ε is a valid plaintext. Naturally, the $\mathcal{B}^{quat}$-coefficients of α, ε' have approximate variance $\sigma_N^2, \sigma_N'^2$ respectively. For C sufficiently large, we have, with overwhelming probability,*

$$||\varepsilon'||_2 < \lambda_1(\mathcal{L}_\alpha)/2 \ \text{ and } \ ||\alpha||_2 < CN^{5/4}\lambda_1(\mathcal{L}_\alpha).$$

Proof. We begin by giving a probable upper bound on $||\varepsilon'||_2$. First note that

$$\mathbb{E}\left(||\varepsilon'||_2\right)^2 < \mathbb{E}\left(||\varepsilon'||_2^2\right) \approx N\sigma_N'^2.$$

Therefore, by Markov's inequality, for any $D > 0$ we have

$$\mathbb{P}\left(||\varepsilon'||_2 < D\sqrt{N}\sigma'_N\right) > 1 - \frac{1}{D}. \tag{11}$$

Similarly, we have

$$\mathbb{P}\left(||\alpha||_2 < D\sqrt{N}\sigma_N\right) > 1 - \frac{1}{D}. \tag{12}$$

It remains to give a probable lower bound on $\lambda_1(\mathcal{L}_\alpha)$. By Minkowski's bound (3), we have

$$\frac{1}{n!}\mathrm{Vol}(\mathcal{L}_\alpha) \leq \lambda_1\lambda_2\cdots\lambda_n. \tag{13}$$

Since $\mathcal{L}_\alpha$ was obtained from the right ideal $\alpha\Lambda$, we find that $\mathcal{L}_\alpha$ is a particularly structured lattice, which will allow us to simplify the above inequality rather dramatically. First note that Λ contains the following distinct norm-one elements:

$$\{\zeta_n^r,\ j\zeta_n^r : 0 \leq r < N\}.$$

Suppose that $\iota(x) \in \mathcal{L}_\alpha$ is a non-zero vector with minimal 2-norm. Since $\alpha\Lambda \leq \Lambda$ is a right ideal, we find that

$$||\iota(x)||_2 = ||\iota(x\zeta_n^r)||_2 = ||\iota(xj\zeta_n^r)||_2,$$

for all $0 \leq r < N$. We deduce that $\alpha\Lambda$ is a well-rounded lattice, that is,

$$\lambda_1 = \lambda_2 = \cdots = \lambda_n.$$

Applying this fact to Minkowski's bound (13), we obtain

$$\frac{1}{n!}\mathrm{Vol}(\mathcal{L}_\alpha) \leq \lambda_1^n. \tag{14}$$

In order to simplify the left-hand side of the above inequality, we will apply Stirling's approximation: $n! \sim \sqrt{2\pi n}(n/e)^n$. We immediately find that

$$\frac{e}{n}|\det(M_\alpha)|^{\frac{1}{n}} = \frac{e}{n}\mathrm{Vol}(\mathcal{L}_\alpha)^{\frac{1}{n}} \lesssim \lambda_1(\mathcal{L}_\alpha).$$

As usual, let $\{\lambda_N^r : 0 \leq r < 2N\}$ denote the eigenvalues of M_α. We have the obvious inequality

$$2eR_{Dec} = \frac{e}{n}\min|\lambda_N^r| \leq \frac{e}{n}|\det(M_\alpha)|^{\frac{1}{n}} \lesssim \lambda_1(\mathcal{L}_\alpha).$$

Recall that we have already given a probable estimate for R_{Dec}. In particular, for C sufficiently large,

$$\lambda_1(\mathcal{L}_\alpha) \gtrsim 2eR_{Dec} > \frac{2D\sigma_N}{CN^{3/4}}, \tag{15}$$

with overwhelming probability. Combining inequalities (11) and (15), we conclude that, for C, D sufficiently large,

$$||\varepsilon'||_2 < D\sqrt{N}\sigma'_N = \frac{D\sigma_N}{CN^{3/4}} < \lambda_1(\mathcal{L}_\alpha)/2$$

with overwhelming probability. On the other hand, for C sufficiently large,

$$||\alpha||_2 \approx \sqrt{N}\sigma_N < CN^{5/4}\lambda_1(\mathcal{L}_\alpha),$$

with overwhelming probability. □

7.2 Security Reductions

We now discuss three aspects of security: one-wayness of encryption, hardness of key-recovery, and IND-CPA security. For ease of exposition, we first define the general setting considered throughout this section:

Setting 1. Suppose that we are given a number of elements $\varepsilon_r \in \Lambda$, each with $||\varepsilon_r||_\infty \leq 1$. We refer to each ε_r as a *plaintext*. Let $\sigma_N > CN^{5/4}$ for some large constant $C > 0$, and define

$$\sigma'_N = \sigma_N / CN^{5/4}.$$

We run the following procedures:

1. **KeyGen**(σ_N). This yields a valid SoliloQuat private key $\alpha \in \Lambda$, where each $\mathcal{B}^{quat}$-coefficient has approximate variance σ_N^2. We denote the corresponding public key as $PK = (\mathfrak{p}, \theta)$.

2. **Enc**$(\varepsilon_r, \sigma'_N, PK)$. During this procedure, we generate a modified plaintext $\varepsilon'_r \in \Lambda$, where each $\mathcal{B}^{quat}$-coefficient has approximate variance σ'^2_N. We denote each corresponding ciphertext as $z_r \in \Lambda$.

Proposition 4 (One-wayness). *Assume that we are in Setting 1, with one plaintext $\varepsilon_1 = \varepsilon$. Suppose that an adversary has access to the ciphertext $z_1 = z$ and the public key PK. The task of recovering $\varepsilon'_1 = \varepsilon'$ is an instance of γ-BDD with the following parameters:*

- *lattice $\mathcal{L} = \mathcal{L}_\alpha$, with basis given by the public representation $\alpha\Lambda = \mathfrak{p}\Lambda + \theta\Lambda$;*
- *radius r, a rational approximation of $\sigma'_N\sqrt{N}$;*
- *target vector $\mathbf{t} = \iota(z)$;*
- *solution vector $\mathbf{v} = \iota(z - \varepsilon')$;*
- *approximation factor $\gamma = O\left(N^{5/4}\right)$.*

Proof. By Proposition 3, for C sufficiently large, we have

$$||\varepsilon'||_2 < \lambda_1(\mathcal{L}_\alpha)/2,$$

with overwhelming probability. Recall equation (10): by construction of the encryption procedure,

$$\iota(z) \equiv \iota(\varepsilon') \pmod{\mathcal{L}_\alpha}.$$

We deduce that $\iota(z - \varepsilon')$ is the closest vector in $\mathcal{L}_\alpha$ to the target vector $\iota(z)$. In order to see that $\gamma = O(N^{5/4})$, we must compare $||\varepsilon'||_2$ and $\lambda_1(\mathcal{L}_\alpha)$. Observe that, for D sufficiently large,

$$\frac{1}{D}\sigma'_N\sqrt{N} < ||\varepsilon'||_2 < \lambda_1(\mathcal{L}_\alpha) < ||\alpha||_2 < D\sigma_N\sqrt{N},$$

with overwhelming probability. Therefore,

$$1 < \frac{\lambda_1(\mathcal{L}_\alpha)}{||\varepsilon'||_2} < O\left(\frac{\sigma_N}{\sigma'_N}\right) = O(N^{5/4}),$$

with overwhelming probability. □

Key Recovery By the construction of SoliloQuat, key recovery can be seen as an instance of the *approximate short generator principal ideal problem:*

Definition 11 (γ-$\mathcal{O}$-SG-PIP). *Suppose that F is a number field, $\mathcal{A}'$ is a quaternion algebra over F, and let $\mathcal{O} \subset \mathcal{A}'$ be an order. Suppose that $\mathcal{I}$ is a principal right $\mathcal{O}$-ideal, and let $\gamma > 1$. Let $\mathbf{B}$ be a $\mathbb{Q}$-basis for $\mathcal{A}'$, and let $||\cdot||_2$ denote the 2-norm on $\mathcal{A}'$ with respect to $\mathbf{B}$, under the coefficient embedding ι. Write $\mathcal{L} := \iota(\mathcal{I})$. Given that there exists a principal generator α of $\mathcal{I}$ with*

$$||\alpha||_2 < \gamma \cdot \lambda_1(\mathcal{L}),$$

find a principal generator $g \in \mathcal{I}$ with $||g||_2 < \gamma \cdot \lambda_1(\mathcal{L})$.

Proposition 5. *Assume that we are in Setting 1. Suppose that an adversary has access to the public key PK. The task of recovering the private key α is an instance of γ-$\mathcal{O}$-SG-PIP with the following parameters:*

- *quaternion algebra $\mathcal{A}' = \mathcal{A}$;*
- *order $\mathcal{O} = \Lambda$;*
- *approximation factor $\gamma < CN^{5/4}$;*
- *base field $F = K$;*
- *right ideal $\mathcal{I} = \alpha\Lambda$ with corresponding lattice $\mathcal{L} = \mathcal{L}_\alpha$;*
- *basis $\mathbf{B} = \mathcal{B}^{quat}$.*

Proof. By Proposition 3, for C sufficiently large, we have

$$||\alpha||_2 < CN^{5/4}\lambda_1(\mathcal{L}_\alpha),$$

with overwhelming probability. The result immediately follows. □

IND-CPA Security We reduce the IND-CPA security of SoliloQuat to decision BDD:

Proposition 6. *Denote the* SoliloQuat *PKE scheme as*

$$\Pi = (\mathbf{KeyGen}, \mathbf{Enc}, \mathbf{Dec}).$$

Consider the IND-CPA experiment $\mathrm{PubK}_{\mathsf{A},\Pi}(n)$*, as given in Definition 10: let the pair of plaintexts selected by the adversary* A *during the IND-CPA experiment be $\varepsilon_1, \varepsilon_2$, and let* **KeyGen** *and* **Enc** *be as in Setting 1 for $r = 1, 2$. Suppose that* A *has advantage Adv in this instance of the IND-CPA experiment. Then* A *has advantage $Adv/2$ in solving decision BDD with the following parameters:*

- *lattice $\mathcal{L} = \mathcal{L}_\alpha$, with basis given by the public representation*

$$\mathcal{L}_\alpha = \iota(\alpha\Lambda) = \iota(\mathfrak{p}\Lambda + \theta\Lambda);$$

- *radius $r := \sigma'_N\sqrt{N}$.*

Proof. We follow a closely related result due to Gentry [12, Theorem 7]. By assumption, A has access to an algorithm $\mathbf{Alg}_1$ that solves the IND-CPA experiment with plaintexts $\varepsilon_1, \varepsilon_2$ with advantage Adv. Suppose that A receives an instance $\mathbf{r}$ of decision BDD parameterised by the lattice $\mathcal{L}$ and radius r; let b denote the bit chosen in this instance of the IND-CPA game. The adversary A can construct an algorithm $\mathbf{Alg}_2$ that solves this instance of BDD with advantage $Adv/2$ as follows: select a random bit $\beta \leftarrow_R \{0,1\}$ and form the ciphertext

$$c \leftarrow \varepsilon_\beta + 2\mathbf{r} \pmod{\mathcal{L}_\alpha}.$$

Suppose that $\mathbf{Alg}_1$ returns the bit β' upon input c. We define the output of $\mathbf{Alg}_2$ as $b' := \beta \oplus \beta'$. If $b = 0$ then c has the same distribution as a valid SoliloQuat encryption of ε_β. Therefore, $\mathbf{Alg}_1$ returns $\beta' = \beta$ with advantage Adv so $\mathbf{Alg}_2$ returns $b' = b$ with advantage Adv. Now suppose that $b = 1$: we first show that there is an isomorphism of additive groups

$$\Lambda/\alpha\Lambda \cong \mathbb{F}_p^2.$$

Indeed, by the proof of Proposition 4, $\Lambda/\mathfrak{p}\Lambda \cong M_2(\mathbb{F}_p)$. Moreover, by [23, Theorem 22.15], $\Lambda/\alpha\Lambda$ is a simple right $\Lambda/\mathfrak{p}\Lambda$-module. It is standard that simple right modules over n-by-n matrix rings over finite fields are isomorphic to $\mathbb{F}_p^n$. We deduce that $\Lambda/\alpha\Lambda \cong \mathbb{F}_p^2$, as claimed. Therefore, $\mathbb{Z}^n/\mathcal{L}_\alpha \cong \mathbb{F}_p^2$. Since p is odd, it follows that $2\mathbf{r}$ is uniformly distributed in $\mathbb{Z}^n/\mathcal{L}_\alpha$. Therefore, c is also uniformly distributed in $\mathbb{Z}^n/\mathcal{L}_\alpha$ so $\mathbf{Alg}_2$ returns $b' = b$ with advantage 0. We conclude that $\mathbf{Alg}_2$ resolves this particular instance of decision BDD with advantage $Adv/2$. □

8 Conclusion

We conclude by revisiting the cautionary tale of Campbell–Groves–Shepherd and by outlining several natural directions for future work on SoliloQuat. The rich structure of maximal orders in cyclotomic number fields can be harnessed to obtain highly efficient and expressive cryptographic constructions, as illustrated by Soliloquy. Yet this same structure is also a poisoned chalice: Campbell–Groves–Shepherd exploited it to break Soliloquy, using the log-unit lattice to reinterpret key recovery as an abelian hidden subgroup problem.

SoliloQuat is motivated by the hope of preserving much of the efficiency and homomorphic functionality of Soliloquy while moving to a setting with fundamentally less accessible structure. Whether this tradeoff is ultimately favorable remains unclear. On one hand, the quaternionic setting may obstruct the specific attacks that succeed in the commutative case; on the other hand, it introduces several new avenues that merit careful study. In particular, we highlight the following questions:

- Do the techniques of Lenstra–Silverberg and Kirchner [20,17] admit meaningful analogues in the quaternionic setting?

– Can key recovery for SoliloQuat be reformulated as an instance of a *non-abelian* hidden subgroup problem?

Beyond cryptanalysis, several questions arise on the constructive side:

– Is the efficiency of SoliloQuat competitive with that of Soliloquy in concrete implementation?
– Can SoliloQuat serve as a basis for fully homomorphic encryption?
– Can the design philosophy behind Soliloquy be extended to other non-commutative algebraic settings, such as the octonions?

These questions are relevant beyond SoliloQuat itself. In particular, since the security of HAWK admits a reduction to the quaternion principal ideal problem [6], any progress on the quaternion *short* principal ideal problem could have consequences for a much wider class of cryptographic constructions. The study of quaternionic ideal problems is not only a niche extension of Soliloquy, but part of a broader effort to understand the constructive promise, and the cryptanalytic risk, of cryptography based on noncommutative algebraic structures.

Acknowledgements The authors thank Joshua Limbrey, Prof. Cong Ling, and Dr. Christian Porter for helpful comments and discussions on this paper. While conducting this research, BN was a research assistant at Imperial College London, under the supervision of Prof. Cong Ling.

References

1. Bernard, O., Roux-Langlois, A.: Twisted-PHS: Using the product formula to solve approx-SVP in ideal lattices. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020. LNCS, vol. 12492, pp. 349–380. Springer International Publishing (2020). https://doi.org/10.1007/978-3-030-64834-3_12
2. Biasse, J.F., Song, F.: Efficient quantum algorithms for computing class groups and solving the principal ideal problem in arbitrary degree number fields. In: SODA 16. p. 893–902. Society for Industrial and Applied Mathematics, USA (2016). https://doi.org/10.5555/2884435.2884499
3. Buchmann, J., Maurer, M., Möller, B.: Cryptography based on number fields with large regulator. Journal de théorie des nombres de Bordeaux **12**(2), 293–307 (2000), https://www.numdam.org/item/JTNB_2000__12_2_293_0/
4. Campbell, P., Groves, M., Shepherd, D.: Soliloquy: A cautionary tale. In: ETSI 2nd Quantum-Safe Crypto Workshop (2014), https://docbox.etsi.org/Workshop/2014/201410_CRYPTO/S07_Systems_and_Attacks/S07_Groves_Annex.pdf
5. Cassels, J.: An Introduction to the Geometry of Numbers. Reprint of the 1971 Edition. Springer (1997). https://doi.org/10.1007/978-3-642-62035-5
6. Chevignard, C., Mureau, G., Espitau, T., Pellet-Mary, A., Pliatsok, H., Wallet, A.: A reduction from Hawk to the principal ideal problem in a quaternion algebra. In: Fehr, S., Fouque, P. (eds.) EUROCRYPT 2025. LNCS, vol. 15602, pp. 154–183. Springer Nature Switzerland (2025). https://doi.org/10.1007/978-3-031-91124-8_6

7. Cramer, R., Ducas, L., Peikert, C., Regev, O.: Recovering short generators of principal ideals in cyclotomic rings. In: Fischlin, M., Coron, J. (eds.) EUROCRYPT 2016. LNCS, vol. 9666, pp. 559–585. Springer Berlin Heidelberg (2016). https://doi.org/10.1007/978-3-662-49896-5_20
8. Cramer, R., Ducas, L., Wesolowski, B.: Short Stickelberger class relations and application to ideal-SVP. In: Coron, J., Nielsen, J.B. (eds.) EUROCRYPT 2017. LNCS, vol. 10210, pp. 324–348. Springer International Publishing (2017). https://doi.org/10.1007/978-3-319-56620-7_12
9. Ducas, L., Plançon, M., Wesolowski, B.: On the shortness of vectors to be found by the ideal-SVP quantum algorithm. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019. LNCS, vol. 11692, pp. 322–351. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26948-7_12
10. Ducas, L., Postlethwaite, E.W., Pulles, L.N., Woerden, W.v.: Hawk: Module LIP makes lattice signatures fast, compact and simple. In: Agrawal, S., Lin, D. (eds.) ASIACRYPT 2022. LNCS, vol. 13794, pp. 65–94. Springer, Cham (2022). https://doi.org/10.1007/978-3-031-22972-5_3
11. Dummit, D.S.: A note on the equivalence of the parity of class numbers and the signature ranks of units in cyclotomic fields (2018), https://arxiv.org/abs/1812.04973
12. Gentry: Fully homomorphic encryption using ideal lattices. In: STOC '09: Symposium on Theory of Computing. p. 169–178. Association for Computing Machinery (2009). https://doi.org/10.1145/1536414.1536440
13. Gentry, C., Szydlo, M.: Cryptanalysis of the revised NTRU signature scheme. In: Knudsen, L.R. (ed.) EUROCRYPT 2002. LNCS, vol. 2332, pp. 299–320. Springer Berlin Heidelberg (2002). https://doi.org/10.1007/3-540-46035-7_20
14. Gentry, C.: A Fully Homomorphic Encryption Scheme. Ph.D. thesis, Stanford University (2009), https://crypto.stanford.edu/craig/
15. Hooley, C.: On Artin's conjecture. Journal für die reine und angewandte Mathematik (1967). https://doi.org/10.1515/crll.1967.225.209
16. Katz, J., Lindell, Y.: Introduction to Modern Cryptography, 2nd Edition. Chapman & Hall/CRC Cryptography and Network Security Series, Taylor & Francis (2014)
17. Kirchner, P.: Algorithms on ideal over complex multiplication order. CoRR **abs/1602.09037** (2016), http://arxiv.org/abs/1602.09037
18. Kirschmer, M., Voight, J.: Algorithmic enumeration of ideal classes for quaternion orders. SIAM J. Comput. **39**, 1714–1747 (2008). https://doi.org/10.5555/1958016.1958018
19. Kocer, E.G.: Circulant, negacyclic and semicirculant matrices with the modified Pell, Jacobsthal and Jacobsthal-Lucas numbers. Hacettepe Journal of Mathematics and Statistics **36**, 133 – 142 (2007), https://izlik.org/JA62CP38CP
20. Lenstra, H.W., Silverberg, A.: Testing isomorphism of lattices over CM-orders. SIAM Journal on Computing **48**(4), 1300–1334 (2019). https://doi.org/10.1137/17M115390X
21. Lyubashevsky, V., Micciancio, D.: Generalized compact knapsacks are collision resistant. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) Automata, Languages and Programming. pp. 144–155. Springer Berlin Heidelberg (2006). https://doi.org/10.1007/11787006_13
22. Pellet-Mary, A., Hanrot, G., Stehlé, D.: Approx-SVP in ideal lattices with preprocessing. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019. LNCS, vol. 11477, pp. 685–716. Springer International Publishing (2019). https://doi.org/10.1007/978-3-030-17656-3_24

23. Reiner, I.: Maximal Orders. London Mathematical Society Monographs, Oxford University Press (2003)
24. Rivest, R.L., Adleman, L., Dertouzos, M.L.: On data banks and privacy homomorphisms. Foundations of Secure Computation pp. 169–179 (1978)
25. Scholz, F.: Central limit theorems and proofs (2019), https://faculty.washington.edu/fscholz/DATAFILES394_2019/CLT.pdf
26. Smart, N.P., Vercauteren, F.: Fully homomorphic encryption with relatively small key and ciphertext sizes. In: Nguyen, P.Q., Pointcheval, D. (eds.) PKC 2010. LNCS, vol. 6056, pp. 420–443. Springer Berlin Heidelberg (2010). https://doi.org/10.1007/978-3-642-13013-7_25
27. The OEIS Foundation Inc.: Entry A005596 in the on-line encyclopedia of integer sequences. https://oeis.org/A005596, accessed 23 February 2026
28. Vershynin, R.: High-Dimensional Probability: An Introduction with Applications in Data Science. Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, Cambridge, 2 edn. (2026)
29. Vignéras, M.F.: Simplification pour les ordres des corps de quaternions totalement définis. Journal für die reine und angewandte Mathematik **0286_0287**, 257–277 (1976), http://eudml.org/doc/151788

A Appendices

A.1 Supplementary Material for Remark 2

We continue with the notation used in Remark 2: suppose that p is a prime congruent to 1 modulo n, and let $w \in \mathbb{Z}$ have order n modulo p. As before, we define $\mathfrak{p} := \mathfrak{p}_1 \cap \mathcal{O}_K$, where $\mathfrak{p}_1 = (p, \zeta_n - w) \leq \mathcal{O}_L$. We are interested in the prime ideals $\mathfrak{q}_1, \mathfrak{q}_2 \leq \mathcal{O}_F$ lying over $\mathfrak{p}$, where

$$F := K(\underbrace{i - r\nu j + rk}_{\theta}) \cong K\left(\sqrt{-1 - r^2(\nu^2+1)}\right).$$

The natural approach is to apply the Dedekind-Kummer theorem; in order to do so, we must first check that $p \nmid [\mathcal{O}_F : \mathcal{O}_K[\theta]]$. Since

$$\mathrm{disc}(\mathcal{O}_K[\theta]) = [\mathcal{O}_F : \mathcal{O}_K[\theta]]^2 \mathrm{disc}(\mathcal{O}_F),$$

it suffices to show that $p \nmid \mathrm{disc}(\mathcal{O}_K[\theta])$. We will verify this with a direct computation. Taking $\{\omega_1, \omega_2, \ldots, \omega_{N/2}\}$ as an integral basis for $\mathcal{O}_K$, we find that $\mathcal{O}_K[\theta]$ has the $\mathbb{Z}$-basis $\{\omega_r, \omega_r\theta : 1 \leq r \leq N/2\}$. Let $\sigma_1, \ldots, \sigma_{N/2}$ denote the embeddings $K \hookrightarrow \mathbb{C}$, and define

$$(M_{rs})_{rs} = \sigma_r(\omega_s),$$

for $1 \leq r, s \leq N/2$. Of course, $\det(M)^2 = \mathrm{disc}(\mathcal{O}_K)$. Similarly, the discriminant of $\mathcal{O}_K[\theta]$ may be computed as

$$\mathrm{disc}(\mathcal{O}_K[\theta]) = \begin{vmatrix} M & M\theta \\ M & -M\theta \end{vmatrix}^2.$$

By elementary row operations, we deduce that

$$\mathrm{disc}(\mathcal{O}_K[\theta]) = \begin{vmatrix} M & M\theta \\ 0 & -2M\theta \end{vmatrix}^2 = \det(M)^4(-2\theta)^N.$$

Since p is unramified in $\mathcal{O}_K$, we know that $p \nmid \det(M)$. On the other hand,

$$(-2\theta)^N = 2^N(-1 - r^2(\nu^2+1))^{N/2}.$$

Since $\nu^2 \equiv -1 \bmod p$, we find that $p \nmid (-2\theta)^N$. We conclude that $p \nmid \mathrm{disc}(\mathcal{O}_K[\theta])$, as required.

We may now apply the Dedekind-Kummer theorem to the extension F/K. The minimal polynomial of θ over K is simply

$$f(X) = X^2 + (1 + r^2(\nu^2+1)).$$

Since $\nu^2 \equiv -1 \bmod p$, the reduction of f modulo p is

$$\overline{f}(X) = X^2 + 1.$$

We deduce that the prime ideals $\mathfrak{q}_1, \mathfrak{q}_2 \leq \mathcal{O}_F$ lying over $\mathfrak{p}$ are precisely

$$\mathfrak{q}_1 = \mathfrak{p}\mathcal{O}_F + (\theta - \nu)\mathcal{O}_F,$$
$$\mathfrak{q}_2 = \mathfrak{p}\mathcal{O}_F + (\theta + \nu)\mathcal{O}_F.$$

We claim that $\mathfrak{q}_2 \cap \mathcal{O}_K[\theta] = (\mathfrak{p}, \theta + \nu)\mathcal{O}_K[\theta]$. From this, it will follow that $\mathfrak{m}_r := \mathfrak{p}\Lambda + (\theta+\nu)\Lambda = (\mathfrak{q}_2 \cap \mathcal{O}_K[\theta])\Lambda$, as required. To see this, we first define

$$R := \mathcal{O}_K[\theta]/\mathfrak{q}_2 \cap \mathcal{O}_K[\theta],$$
$$S := \mathcal{O}_K[\theta]/(\mathfrak{p}, \theta+\nu)\mathcal{O}_K[\theta].$$

Since $\mathfrak{q}_2 \cap \mathcal{O}_K[\theta] \supset (\mathfrak{p}, \theta+\nu)\mathcal{O}_K[\theta]$, R is a quotient of S. Moreover, since $\mathfrak{p}$ is completely split in $\mathcal{O}_K$, we know that $|S| = p$. Therefore, we either have $|R| = 1$ or $R \cong S$. Clearly, $\mathfrak{q}_2 \cap \mathcal{O}_K[\theta] \neq \mathcal{O}_K[\theta]$, so we must have $R \cong S$. We deduce that $|R| = |S|$, so in fact $(\mathfrak{p}, \theta+\nu)\mathcal{O}_K[\theta] = \mathfrak{q}_2 \cap \mathcal{O}_K[\theta]$, as we wanted to show.

A.2 Testing Principality

Let $\mathfrak{m}_r \leq \Lambda$ be a maximal right ideal lying over $\mathfrak{p} = (p, \zeta_n - w) \cap \mathcal{O}_K$, where $w \in \mathbb{Z}$ has order n modulo p. We will assume that $\mathfrak{m}_r$ has a known principal generator $\alpha \in \Lambda$ with

$$\mathrm{nrd}(\alpha)\mathcal{O}_K = \mathfrak{p}.$$

Define $\nu = w^{n/4}$ and consider the number field

$$F = K(\underbrace{i - r\nu j + rk}_{\theta}) \cong K\left(\sqrt{-1 - r^2(\nu^2+1)}\right).$$

In Appendix A.1, we demonstrated that $\mathfrak{m}_r = (\mathfrak{q} \cap \mathcal{O}_K[\theta])\Lambda$, for some $\mathfrak{q} \leq \mathcal{O}_F$ lying over $\mathfrak{p}$. In this section, we describe a deterministic polynomial time algorithm that tests if $\mathfrak{q} \cap \mathcal{O}_K[\theta] \leq \mathcal{O}_K[\theta]$ is principal. For sake of argument, assume that $\mathfrak{q} \cap \mathcal{O}_K[\theta] = g\mathcal{O}_K[\theta]$ for some $g \in \mathcal{O}_K[\theta]$. Since $\mathfrak{m}_r = g\Lambda$, we find that

$$\mathrm{nrd}(\alpha)\mathcal{O}_K = \mathrm{nrd}(\alpha\Lambda) = \mathrm{nrd}(g\Lambda) = \mathrm{nrd}(g)\mathcal{O}_K.$$

Therefore, $\mathrm{nrd}(\alpha) = \mathrm{nrd}(g)u$ for some $u \in \mathcal{O}_K^\times$. Since $\mathrm{nrd}(\alpha)$ and $\mathrm{nrd}(g)$ are both totally positive, u must also be totally positive. As in [11, Proposition 1, statement 5.b2], the totally positive units of $\mathcal{O}_K$ are precisely the squares of units in $\mathcal{O}_K$. Therefore, there exists $v \in \mathcal{O}_K^\times$ such that $u = v^2$. Putting this all together, we see that

$$\mathrm{nrd}(\alpha) = \mathrm{nrd}(gv).$$

Moreover, gv is a principal generator of $\mathfrak{q} \cap \mathcal{O}_K[\theta]$ since $v \in \mathcal{O}_K^\times$.

From now on, we no longer assume that $\mathfrak{q} \cap \mathcal{O}_K[\theta]$ is principal. By [20, Theorem 1.3], there is a deterministic polynomial time algorithm that, given $\mathrm{nrd}(\alpha)$, decides if there exists a principal generator g of $\mathfrak{q} \cap \mathcal{O}_K[\theta]$ with $g\overline{g} = \mathrm{nrd}(\alpha)$. As above, there exists such a principal generator g if, and only if, $\mathfrak{q} \cap \mathcal{O}_K[\theta]$ is principal. Therefore, there is a deterministic polynomial time algorithm that decides whether $\mathfrak{q} \cap \mathcal{O}_K[\theta]$ is principal or not.

A.3 Supplementary Material for Section 6

Lemma 1. Suppose that $\beta, \varepsilon' \in \mathcal{A}$. Then

$$||\beta\varepsilon'||_\infty \leq 2N||\beta||_\infty||\varepsilon'||_\infty.$$

Proof. We begin by defining the following quantities:

$$\delta_\infty = \sup\left\{\frac{||cd||_\infty}{||c||_\infty||d||_\infty} : c, d \in \mathcal{O}_L\right\},$$
$$\delta_\infty^{quat} = \sup\left\{\frac{||cd||_\infty}{||c||_\infty||d||_\infty} : c, d \in \Lambda\right\}.$$

The δ_∞ constant was first introduced by Lyubashevsky and Micciancio [21], and was later studied by Gentry [14]. In this work (his thesis), Gentry derives several upper bounds on the quantity. This definition will prove useful, allowing us to relate $||\beta\varepsilon'||_\infty$ to $||\beta||_\infty||\varepsilon'||_\infty$. Suppose that we are given $c, d \in \Lambda$. Of course, we may write

$$c = c_1 + jc_2,$$
$$d = d_1 + jd_2,$$

where $c_1, c_2, d_1, d_2 \in \mathcal{O}_L$. After a straightforward computation, we find that

$$cd = (c_1d_1 - \overline{c_2}d_2) + j(\overline{c_1}d_2 + c_2d_1). \tag{16}$$

Note that, for $x, y \in \mathcal{O}_L$, we have $||x + jy||_\infty = \max(||x||_\infty, ||y||_\infty)$. Applying this to equation (16), we find that

$$||cd||_\infty \leq \max(||c_1 d_1 - \overline{c_2} d_2||_\infty, ||\overline{c_1} d_2 + c_2 d_1||_\infty). \tag{17}$$

By the triangle inequality,

$$\begin{aligned} ||c_1 d_1 - \overline{c_2} d_2||_\infty &\leq ||c_1 d_1||_\infty + ||\overline{c_2} d_2||_\infty \\ &\leq \delta_\infty(||c_1||_\infty ||d_1||_\infty + ||\overline{c_2}||_\infty ||d_2||_\infty) \\ &\leq 2\delta_\infty ||c||_\infty ||d||_\infty. \end{aligned}$$

Similarly, $||\overline{c_1} d_2 + c_2 d_1||_\infty \leq 2\delta_\infty ||c||_\infty ||d||_\infty$. Returning to inequality (17), we deduce that

$$\frac{||cd||_\infty}{||c||_\infty ||d||_\infty} \leq 2\delta_\infty.$$

In other words, $\delta_\infty^{quat} \leq 2\delta_\infty$. We can't quite use δ_∞^{quat} to relate $||\beta\varepsilon'||_\infty$ to $||\beta||_\infty ||\varepsilon'||_\infty$; we defined δ_∞^{quat} as a supremum over Λ, while β is an element of $\mathcal{A}$. This distinction does not matter. Indeed, note that there exists an integer m such that $m\beta \in \Lambda$. For example, one may take m to be the product of the denominators of each coefficient $b_r^{(1)}, b_r^{(2)} \in \mathbb{Q}$. By the definition of δ_∞^{quat}, we have

$$m||\beta\varepsilon'||_\infty = ||m\beta\varepsilon'||_\infty \leq \delta_\infty^{quat} ||m\beta||_\infty ||\varepsilon'||_\infty \leq 2m\delta_\infty ||\beta||_\infty ||\varepsilon'||_\infty.$$

It follows that

$$||\beta\varepsilon'||_\infty \leq 2\delta_\infty ||\beta||_\infty ||\varepsilon'||_\infty.$$

From [26, Lemma 2] we know that $\delta_\infty \leq N$. Applying this fact to the above inequality concludes the proof. □

Before we give a proof of Lemma 2, we further study the left-multiplication map $L_\alpha : x \mapsto \alpha x$. Write

$$\alpha_1 = \sum_{r=0}^{N-1} a_r^{(1)} \zeta_n^r \quad \text{and} \quad \alpha_2 = \sum_{r=0}^{N-1} a_r^{(2)} \zeta_n^r,$$

so that $\alpha = \alpha_1 + j\alpha_2$. Recall that we take the natural $\mathbb{Q}$-basis of $\mathcal{A}$ to be

$$\mathcal{B}^{quat} = \{1, \zeta_n, \ldots, \zeta_n^{N-1}, j, j\zeta_n, \ldots, j\zeta_n^{N-1}\}.$$

Observe the following multiplication rules:

$$\begin{aligned} \alpha\zeta_n^r &= (\alpha_1 + j\alpha_2)\zeta_n^r = \alpha_1 \zeta_n^r + j\alpha_2 \zeta_n^r, \\ \alpha j\zeta_n^r &= (\alpha_1 + j\alpha_2) j\zeta_n^r = -\overline{\alpha_2}\zeta_n^r + j\overline{\alpha_1}\zeta_n^r. \end{aligned}$$

Extending $\mathbb{Q}$-linearly, we find that for $x, y \in L$,

$$\begin{aligned} \alpha x &= \alpha_1 x + j\alpha_2 x \\ \alpha j y &= -\overline{\alpha_2} y + j\overline{\alpha_1} y. \end{aligned}$$

These relations allow us to formulate multiplication in $\mathcal{A}$ as a matrix equation. Suppose that $x + jy \in \mathcal{A}$ where $x, y \in L$. We find that

$$\begin{pmatrix} \alpha_1 & -\overline{\alpha_2} \\ \alpha_2 & \overline{\alpha_1} \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} \alpha_1 x - \overline{\alpha_2} y \\ \alpha_2 x + \overline{\alpha_1} y \end{pmatrix} = \begin{pmatrix} \Re_j(\alpha(x+jy)) \\ \Im_j(\alpha(x+jy)) \end{pmatrix}, \tag{18}$$

where $\Re_j(a+jb) = a$ and $\Im_j(a+jb) = b$ for $a, b \in L$. Note that $\Im_j(a+jb) \neq \Im_j(a+bj)$ in general. We can also formulate multiplication in L as a matrix equation. Let A_r denote the representation of the map $L \to L$ given by $x \mapsto \alpha_r x$ with respect to $\mathcal{B}$. Explicitly, we have

$$A_r := \begin{pmatrix} a_0^{(r)} & -a_{N-1}^{(r)} & -a_{N-2}^{(r)} & \cdots & -a_1^{(r)} \\ a_1^{(r)} & a_0^{(r)} & -a_{N-1}^{(r)} & \cdots & -a_2^{(r)} \\ a_2^{(r)} & a_1^{(r)} & a_0^{(r)} & \cdots & -a_3^{(r)} \\ \vdots & \vdots & \vdots & & \vdots \\ a_{N-1}^{(r)} & a_{N-2}^{(r)} & a_{N-3}^{(r)} & \cdots & a_0^{(r)} \end{pmatrix}.$$

Note that the map $L \to L$ given by $x \mapsto \overline{\alpha_r} x$ is represented by $A_r^\intercal$ with respect to $\mathcal{B}$. This can be seen from the fact that

$$\overline{\alpha_1} = \sum_{r=0}^{N-1} a_r^{(1)} \zeta_n^{-r} = a_0^{(1)} - \sum_{r=1}^{N-1} a_r^{(1)} \zeta_n^{N-r} = a_0^{(1)} - \sum_{r=1}^{N-1} a_{N-r}^{(1)} \zeta_n^r.$$

Therefore, the coefficients of $\overline{\alpha_r}$ with respect to $\mathcal{B}$ are precisely those occurring in the first column of $A_r^\intercal$. Thus, the map $x \mapsto \overline{\alpha_r} x$ is indeed represented by $A_r^\intercal$ with respect to $\mathcal{B}$. Replacing $\alpha_r \leftrightarrow A_r$ in equation (18), we find that L_α is represented by the matrix

$$M_\alpha = \begin{pmatrix} A_1 & -A_2^\intercal \\ A_2 & A_1^\intercal \end{pmatrix}. \tag{19}$$

Lemma 2. Suppose we are given an element $0 \neq \alpha \in \Lambda$ and define $\beta = \alpha^{-1}$. Let M_α have eigenvalues $\lambda_N^0, \lambda_N^1, \ldots, \lambda_N^{2N-1}$. Then

$$||\beta||_\infty \leq \frac{1}{\min_r |\lambda_N^r|}.$$

Proof. Since $\mathcal{B}$ is a $\mathbb{Q}$-basis for $\mathcal{A}$, and $\mathbb{Q} \subseteq \mathcal{A}$ is central, we have that

$$M_\alpha M_\beta = M_{\alpha\beta} = M_1 = I_{2N}.$$

Note that the first column of M_β is given by

$$\mathbf{b} = \left(b_0^{(1)} \; b_1^{(1)} \cdots b_{N-1}^{(1)} \; b_0^{(2)} \; b_1^{(2)} \cdots b_{N-1}^{(2)} \right)^\intercal.$$

We deduce that $M_\alpha \mathbf{b} = \mathbf{e}_1$. Let M_α^r denote the matrix formed by replacing the $(r+1)^{th}$ column of M_α with the vector $\mathbf{e}_1$. By Cramer's rule, we have

$$b_r^{(1)} = \frac{\det(M_\alpha^r)}{\det(M_\alpha)}, \; 0 \leq r \leq N-1,$$
$$b_{r-N}^{(2)} = \frac{\det(M_\alpha^r)}{\det(M_\alpha)}, \; N \leq r \leq 2N-1.$$

We therefore have the uniform bound

$$||\mathbf{b}||_\infty \leq \frac{\max_r |\det(M_\alpha^r)|}{|\det(M_\alpha)|}. \tag{20}$$

Recall that the 2-norm $||\mathbf{A}||_2$ of a matrix $\mathbf{A} \in M_{2N}(\mathbb{C})$ is defined as

$$||\mathbf{A}||_2 = \max_{\mathbf{x}\neq 0} \frac{||\mathbf{A}\mathbf{x}||_2}{||\mathbf{x}||_2}.$$

From this definition, for any given r, we have the inequality

$$|\det(M_\alpha^r)| \leq ||\mathrm{Adj}(M_\alpha)||_2.$$

The 2-norm can alternatively be defined as

$$||\mathbf{A}||_2 = \sqrt{\rho(\mathbf{A}^*\mathbf{A})},$$

where $\mathbf{A}^*$ denotes the conjugate transpose of $\mathbf{A}$, and $\rho(\mathbf{A}^*\mathbf{A})$ denotes the maximum absolute value of an eigenvalue of $\mathbf{A}^*\mathbf{A}$, which is known as the *spectral radius*. Taking $\mathbf{A} = \mathrm{Adj}(M_\alpha)$, we find that

$$|\det(M_\alpha^r)| \leq \sqrt{\rho(\mathrm{Adj}(M_\alpha)^*\mathrm{Adj}(M_\alpha))}. \tag{21}$$

Since M_α is normal, so is the adjugate $\mathrm{Adj}(M_\alpha)$. It follows that

$$\sqrt{\rho(\mathrm{Adj}(M_\alpha)^*\mathrm{Adj}(M_\alpha))} = \rho(\mathrm{Adj}(M_\alpha)).$$

This allows us to rephrase inequality (21) as

$$|\det(M_\alpha^r)| \leq \rho(\mathrm{Adj}(M_\alpha)), \tag{22}$$

for any r. Using the fact that $\mathrm{Adj}(M_\alpha) = \det(M_\alpha)M_\alpha^{-1}$, we find that the eigenvalues of $\mathrm{Adj}(M_\alpha)$ are precisely

$$\left\{ \prod_{r\neq s} \lambda_N^r : 0 \leq s \leq 2N-1 \right\}.$$

Since M_α is non-singular, we therefore have

$$\rho(\mathrm{Adj}(M_\alpha)) = \frac{|\det(M_\alpha)|}{\min_r |\lambda_N^r|}. \tag{23}$$

Combining equation (23) and inequality (22), we obtain the upper bound, for all r,

$$|\det(M_\alpha^r)| \leq \frac{|\det(M_\alpha)|}{\min_r |\lambda_N^r|}.$$

Applying this to inequality (20), we obtain the final bound

$$||\beta||_\infty = ||\mathbf{b}||_\infty \leq \frac{\frac{|\det(M_\alpha)|}{\min_r |\lambda_N^r|}}{|\det(M_\alpha)|} = \frac{1}{\min_r |\lambda_N^r|}. \qquad \square$$

A.4 Distribution of Eigenvalues

Let $a_r^{(1)}, a_r^{(2)}$ be IID random variables with distribution $\mathcal{D}_{\mathbb{Z},\sigma_N}$. We can form the related random variable

$$\alpha = \sum_{r=0}^{N-1} a_r^{(1)} \zeta_n^r + j \sum_{r=0}^{N-1} a_r^{(2)} \zeta_n^r.$$

We begin by describing the eigenvalues of M_α. Since M_α is composed of four negacyclic blocks, it will prove useful to understand the eigenvalues of a general negacyclic matrix. Given a vector $\mathbf{x} \in \mathbb{C}^n$, write $N(\mathbf{x})$ for the negacyclic matrix with first column $\mathbf{x}$.

Theorem 2. *There exists a unitary matrix G such that, for any $m \times m$ negacyclic matrix $N(\mathbf{x})$,*

$$G^* N(\mathbf{x}) G = \operatorname{diag}(\mu_0(\mathbf{x}), \ldots, \mu_{m-1}(\mathbf{x})),$$

where

$$\mu_s(\mathbf{x}) = \sum_{r=0}^{m-1} x_r \zeta_m^{\frac{(2s+1)r}{2}}.$$

Proof. This is stated and proven in [19, Theorem 1.2]. □

Recall that

$$M_\alpha = \begin{pmatrix} A_1 & -A_2^\intercal \\ A_2 & A_1^\intercal \end{pmatrix},$$

where A_1, A_2 are negacyclic. Let G denote the matrix given in Theorem 2. Then

$$\begin{pmatrix} G^* & 0 \\ 0 & G^* \end{pmatrix} \begin{pmatrix} A_1 & -A_2^\intercal \\ A_2 & A_1^\intercal \end{pmatrix} \begin{pmatrix} G & 0 \\ 0 & G \end{pmatrix} = \begin{pmatrix} G^* A_1 G & -G^* A_2^\intercal G \\ G^* A_2 G & G^* A_1^\intercal G \end{pmatrix} = \begin{pmatrix} D_1 & -D_2^* \\ D_2 & D_1^* \end{pmatrix},$$

where

$$D_1 = \operatorname{diag}(\lambda_{N,0}^{(1)}, \lambda_{N,1}^{(1)}, \ldots, \lambda_{N,N-1}^{(1)}),$$
$$D_2 = \operatorname{diag}(\lambda_{N,0}^{(2)}, \lambda_{N,1}^{(2)}, \ldots, \lambda_{N,N-1}^{(2)}),$$

and

$$\lambda_{N,s}^{(1)} = \sum_{r=0}^{N-1} a_r^{(1)} \zeta_N^{\frac{(2s+1)r}{2}},$$
$$\lambda_{N,s}^{(2)} = \sum_{r=0}^{N-1} a_r^{(2)} \zeta_N^{\frac{(2s+1)r}{2}}.$$

Suppose that $\sigma \in S_{2N}$ and let $\mathbf{M}$ be an arbitrary $(2N) \times (2N)$ matrix. Recall that, after conjugating by a suitable sequence of elementary matrices, we can obtain a new matrix $\mathbf{M}'$ with entries

$$\mathbf{M}'_{rs} = \mathbf{M}_{\sigma(r)\sigma(s)}.$$

Now take $\sigma \in S_{2N}$ to be the "perfect shuffle", that is,

$$\sigma^{-1}(r) = \begin{cases} 2r-1 & 1 \le r \le N, \\ 2(r-N) & N+1 \le r \le 2N. \end{cases}.$$

By the above discussion, there is an invertible matrix E such that

$$E\mathbf{M}E^{-1} = (\mathbf{M}_{\sigma(r)\sigma(s)})_{r,s}.$$

In particular, we find that

$$E\begin{pmatrix} D_1 & -D_2^* \\ D_2 & D_1^* \end{pmatrix}E^{-1} = \begin{pmatrix} B_0 & & & \\ & B_1 & & \\ & & \ddots & \\ & & & B_{N-1} \end{pmatrix},$$

where each B_s is the 2×2 matrix given by

$$\begin{pmatrix} \lambda_{N,s}^{(1)} & -\overline{\lambda_{N,s}^{(2)}} \\ \lambda_{N,s}^{(2)} & \overline{\lambda_{N,s}^{(1)}} \end{pmatrix}.$$

Note that the characteristic polynomial of B_s is

$$X^2 - \mathrm{tr}(B_s)X + \det(B_s) = X^2 - 2\Re(\lambda_{N,s}^{(1)})X + (|\lambda_{N,s}^{(1)}|^2 + |\lambda_{N,s}^{(2)}|^2).$$

Therefore, the eigenvalues of B_s are given by

$$\Re(\lambda_{N,s}^{(1)}) \pm i\sqrt{(|\lambda_{N,s}^{(1)}|^2 + |\lambda_{N,s}^{(2)}|^2) - (\Re(\lambda_{N,s}^{(1)}))^2}.$$

It follows that each eigenvalue of B_s has absolute value

$$\sqrt{|\lambda_{N,s}^{(1)}|^2 + |\lambda_{N,s}^{(2)}|^2}.$$

Recall that the eigenvalues of a matrix, and their multiplicities, are invariant under conjugation. Therefore, the eigenvalues of M_α have absolute values in the set

$$\left\{ \sqrt{|\lambda_{N,s}^{(1)}|^2 + |\lambda_{N,s}^{(2)}|^2} : 0 \le s \le N-1 \right\},$$

and each absolute value occurs with multiplicity 2. Thus, in order to understand the radius of correct decryption, we need only understand (with high probability) the smallest value occurring in the above set. In order to achieve this, we will use a variant of the central limit theorem due to Lindeberg:

Theorem 3 (Lindeberg's conditions). *A triangular array is a double-indexed family $\{\xi_{r,s}\}_{1\le r,1\le s\le m(r)}$ of random variables. We say that a triangular array $\{\xi_{r,s}\}_{1\le r,1\le s\le m(r)}$ of independent random variables satisfies Lindeberg's conditions if*

1. $\mathbb{E}(\xi_{r,s}) = 0$ *for all* r, s.
2. *For any* n, $\sum_{1 \le s \le m(r)} \mathbb{E}(\xi_{r,s}^2) = 1$.
3. *For every* $\delta > 0$,
$$\lim_{r \to \infty} \sum_{s=1}^{m(r)} \mathbb{E}(\xi_{r,s}^2 \mathbf{1}_{|\xi_{r,s}| \ge \delta}) = 0.$$

Any triangular array $\{\xi_{r,s}\}$ *that satisfies Lindeberg's conditions has, as* $r \to \infty$,

$$\sum_{s=1}^{m(r)} \xi_{r,s} \xrightarrow{d} \mathcal{N}(0,1).$$

Proof of Lindeberg's conditions. See [25]. □

The following proposition describes the asymptotic distribution of the eigenvalues $\lambda_{N,s}^{(1)}, \lambda_{N,s}^{(2)}$:

Proposition 7. *Let* σ_N *be such that* $\sigma_N \to \infty$ *as* $N \to \infty$. *Let* $a_{r,N}^{(1)}$ *be IID random variables with distribution* $\mathcal{D}_{\mathbb{Z},\sigma_N}$. *For each* $0 \le s < N$, *define the random variable*

$$\lambda_{N,s}^{(1)} = \sum_{r=0}^{N-1} a_{r,N}^{(1)} \zeta_n^{(2s+1)r}.$$

For any given s, *we have the following convergence in distribution*

$$\frac{1}{N\sigma_N^2} |\lambda_{N,s}^{(1)}|^2 \xrightarrow{d} \chi_2^2,$$

as $N \to \infty$.

Remark 5. We note that Proposition 7 isn't surprising. Indeed, if we were to instead take $a_{r,N}^{(1)}$ to be IID *continuous* Gaussians, then the result would be trivial. On the other hand, since $\sigma_N \to \infty$, the width of each $a_{r,N}^{(1)}$ will (asymptotically) exceed the smoothing parameter. Therefore, we can expect the discrete case to closely resemble the continuous case.

Before we start the formal proof, we give a brief outline as follows:

1. First, we show that, for any $a, b \in \mathbb{R}$,
$$\frac{a}{\sigma_N \sqrt{N}} \Re(\lambda_{N,s}^{(1)}) + \frac{b}{\sigma_N \sqrt{N}} \Im(\lambda_{N,s}^{(1)}) \xrightarrow{d} \mathcal{N}(0, a^2 + b^2).$$
2. Informally, it follows that $\Re(\lambda_{N,s}^{(1)})/\sigma_N\sqrt{N}$ and $\Im(\lambda_{N,s}^{(1)})/\sigma_N\sqrt{N}$ asymptotically behave as IID standard Gaussians.

3. Using the usual formula

$$|\lambda_{N,s}^{(1)}|^2 = \Re(\lambda_{N,s})^2 + \Im(\lambda_{N,s})^2,$$

we can then deduce the result.

Proof of Proposition 7. As promised, we first show that, for any $a, b \in \mathbb{R}$,

$$\frac{a}{\sigma_N\sqrt{N}}\Re(\lambda_{N,s}^{(1)}) + \frac{b}{\sigma_N\sqrt{N}}\Im(\lambda_{N,s}^{(1)}) \xrightarrow{d} \mathcal{N}(0, a^2+b^2). \tag{24}$$

From the definition of $\lambda_{N,s}^{(1)}$, we have the explicit expressions

$$\Re(\lambda_{N,s}^{(1)}) = \sum_{r=0}^{N-1} a_{r,N}^{(1)} \cos\left(\frac{2\pi(2s+1)r}{n}\right), \tag{25}$$

$$\Im(\lambda_{N,s}^{(1)}) = \sum_{r=0}^{N-1} a_{r,N}^{(1)} \sin\left(\frac{2\pi(2s+1)r}{n}\right). \tag{26}$$

Therefore,

$$\begin{aligned}&\frac{a}{\sigma_N\sqrt{N}}\Re(\lambda_{N,s}^{(1)}) + \frac{b}{\sigma_N\sqrt{N}}\Im(\lambda_{N,s}^{(1)})\\ &\quad= \sum_{r=0}^{N-1}\left(\frac{a}{\sigma_N\sqrt{N}}\cos\left(\frac{2\pi(2s+1)r}{n}\right) + \frac{b}{\sigma_N\sqrt{N}}\sin\left(\frac{2\pi(2s+1)r}{n}\right)\right)a_{r,N}^{(1)}.\end{aligned}$$

Fix s and consider the triangular array $\{\xi_{N,r}\}_{1\le N, 0\le r<N}$ given by

$$\xi_{N,r} = \underbrace{\left(\frac{a}{\sigma_N\sqrt{N}}\cos\left(\frac{2\pi(2s+1)r}{n}\right) + \frac{b}{\sigma_N\sqrt{N}}\sin\left(\frac{2\pi(2s+1)r}{n}\right)\right)}_{C_{N,r,s}} a_{r,N}^{(1)}.$$

By construction,

$$\frac{a}{\sigma_N\sqrt{N}}\Re(\lambda_{N,s}^{(1)}) + \frac{b}{\sigma_N\sqrt{N}}\Im(\lambda_{N,s}^{(1)}) = \sum_{r=0}^{N-1}\xi_{N,r}.$$

From now on, assume that $a^2+b^2=1$. In order to form a triangular array for which Lindeberg's second condition holds, we introduce the constants

$$\gamma_N^2 = \sum_{0\le r<N}\mathbb{E}(\xi_{N,r}^2).$$

Then, if we can verify Lindeberg's conditions for the triangular array

$$\left\{\frac{1}{\gamma_N}\xi_{N,r}\right\}_{1\le N, 0\le r<N},$$

then it will follow that

$$\frac{a}{\gamma_N \sigma_N \sqrt{N}} \Re(\lambda_{N,s}^{(1)}) + \frac{b}{\gamma_N \sigma_N \sqrt{N}} \Im(\lambda_{N,s}^{(1)}) \xrightarrow{d} \mathcal{N}(0,1).$$

Note that we have already claimed that the above expression holds *without* the constants γ_N. In order to show that these claims do not contradict each other, we briefly explain why $\gamma_N \to 1$ as $N \to \infty$. Since $\sigma_N \to \infty$ as $N \to \infty$, we have the (good) approximation $\mathbb{E}((a_i^{(1)})^2) \approx \sigma_N^2$. Therefore, γ_N^2 is approximately

$$\begin{aligned}
&\frac{1}{N} \sum_{r=0}^{N-1} \left(a^2 \cos^2 \left(\frac{2\pi(2s+1)r}{n} \right) + b^2 \sin^2 \left(\frac{2\pi(2s+1)r}{n} \right) + ab \sin \left(\frac{2\pi(2s+1)r}{N} \right) \right) \\
&= \frac{a^2}{N} \sum_{r=0}^{N-1} \cos^2 \left(\frac{2\pi(2s+1)r}{n} \right) + \frac{b^2}{N} \sum_{r=0}^{N-1} \sin^2 \left(\frac{2\pi(2s+1)r}{n} \right) + \frac{ab}{N} \sum_{r=0}^{N-1} \sin \left(\frac{2\pi r}{N} \right) \\
&= \frac{a^2}{N} \left(\frac{N}{2} + O(1) \right) + \frac{b^2}{N} \left(\frac{N}{2} + O(1) \right) \\
&= 1 + O(1/N).
\end{aligned}$$

In other words, $\gamma_N^2 = 1 + O(1/N)$, so we certainly have $\gamma_N \to 1$ as $N \to \infty$. We now show that Lindeberg's conditions do in fact hold for the triangular array

$$\{\xi_{N,r}/\gamma_N\}_{1 \le N, 0 \le r < N}.$$

Lindeberg's 1st *condition:* For all $1 \le N$ and $0 \le r < N$, we have

$$\mathbb{E}\left(\frac{1}{\gamma_N} \xi_{N,r} \right) = \frac{C_{N,r,s}}{\gamma_N} \mathbb{E}(a_{r,N}^{(1)}) = 0.$$

Lindeberg's 2nd *condition:* For all $1 \le N$ and $0 \le r < N$, we have

$$\sum_{0 \le r < N} \mathbb{E}\left(\frac{1}{\gamma_N^2} \xi_{N,r}^2 \right) = \frac{\sum_{0 \le r < N} \mathbb{E}(\xi_{N,r}^2)}{\sum_{0 \le r < N} \mathbb{E}(\xi_{N,r}^2)} = 1.$$

Lindeberg's 3rd *condition:* Given $\delta > 0$, we wish to show that

$$\lim_{N \to \infty} \frac{1}{\gamma_N^2} \sum_{r=0}^{N-1} \mathbb{E}((\xi_{N,r})^2 \mathbf{1}_{|\xi_{N,r}| \ge \delta}) = 0.$$

Since $\gamma_N \to 1$ as $N \to \infty$, we can equivalently show that

$$\lim_{N \to \infty} \sum_{r=0}^{N-1} \mathbb{E}((\xi_{N,r})^2 \mathbf{1}_{|\xi_{N,r}| \ge \delta}) = 0. \tag{27}$$

Substituting the formula for $\xi_{N,r}$ into equation (27), Lindeberg's 3rd condition becomes

$$\lim_{N \to \infty} \sum_{r=0}^{N-1} \mathbb{E}\left(C_{N,r,s}^2 \left(a_{r,N}^{(1)} \right)^2 \mathbf{1}_{|C_{N,r,s} a_{r,N}^{(1)}| \ge \delta} \right) = 0. \tag{28}$$

Note that we have the upper bound

$$C_{N,r,s}^2 \leq \frac{1}{N\sigma_N^2}(|a|+|b|)^2 \leq \frac{4}{N\sigma_N^2}. \tag{29}$$

Combining equation (28) and inequality (29), the claim becomes

$$\begin{aligned}\lim_{N\to\infty} \sum_{r=0}^{N-1} \mathbb{E}\left(\frac{2}{N\sigma_N^2}(a_{r,N}^{(1)})^2 \mathbf{1}_{\frac{2}{\sigma_N\sqrt{N}}|a_{r,N}^{(1)}|\geq\delta}\right) \\ = \lim_{N\to\infty} \frac{2}{\sigma_N^2}\mathbb{E}\left((a_{0,N}^{(1)})^2 \mathbf{1}_{\frac{2}{\sigma_N\sqrt{N}}|a_{0,N}^{(1)}|\geq\delta}\right) = 0.\end{aligned} \tag{30}$$

Since $a_{0,N}^{(1)} \sim \mathcal{D}_{\mathbb{Z},\sigma_N}$, by definition, we have

$$\mathbb{P}\left(a_{0,N}^{(1)} = c\right) = \frac{e^{-\frac{c^2}{2\sigma_N^2}}}{\sum_{l\in\mathbb{Z}} e^{-\frac{l^2}{2\sigma_N^2}}}.$$

Therefore,

$$\frac{1}{\sigma_N^2}\mathbb{E}\left((a_{0,N}^{(1)})^2 \mathbf{1}_{|a_{0,N}^{(1)}|\geq(\delta/2)\sigma_N\sqrt{N}}\right) = \frac{2}{\sigma_N^2} \sum_{r=\lceil(\delta/2)\sigma_N\sqrt{N}\rceil}^{\infty} r^2 \frac{e^{-\frac{r^2}{2\sigma_N^2}}}{\sum_{l\in\mathbb{Z}} e^{-\frac{l^2}{2\sigma_N^2}}}. \tag{31}$$

We will now give a long sequence of upper bounds on the quantity, in order to finally show that the limit as $N\to\infty$, is zero. Since $\delta > 0$ was arbitrary, we will now replace $\delta/2$ with δ for notational simplicity. Note that

$$\frac{1}{\sigma_N}\sum_{l\in\mathbb{Z}} e^{-\frac{(l/\sigma_N)^2}{2}} = \sqrt{2\pi} + O\left(\frac{1}{\sigma_N}\right). \tag{32}$$

This identity comes from approximating the integral

$$\int_{-\infty}^{\infty} e^{-\frac{x^2}{2}}\,\mathrm{d}x$$

by an infinite sum of rectangular areas of width $1/\sigma_N$. The (dominating) error in this approximation is at most

$$\frac{1}{\sigma_N}\sum_{r\in\mathbb{Z}}\left|\left(-\frac{r}{\sigma_N}\right)e^{-\frac{r^2}{2\sigma_N^2}}\right| \approx \frac{2}{\sigma_N}\left(\int_0^{\infty} xe^{-\frac{x^2}{2}}\,\mathrm{d}x\right) = \frac{2}{\sigma_N}\left[-e^{-\frac{x^2}{2}}\right]_0^{\infty} = \frac{2}{\sigma_N},$$

providing the $O(1/\sigma_N)$ term in (32). Using identity (32), quantity (31) becomes

$$\begin{aligned}\frac{1}{\sigma_N^2}\mathbb{E}\left((a_{0,N}^{(1)})^2 \mathbf{1}_{|a_{0,N}^{(1)}|\geq\delta\sigma_N\sqrt{N}}\right) \\ = \frac{2}{\sigma_N\sqrt{2\pi}+O(1)} \sum_{r=\lceil\delta\sigma_N\sqrt{N}\rceil}^{\infty} (r/\sigma_N)^2 e^{-\frac{r^2}{2\sigma_N^2}}.\end{aligned} \tag{33}$$

Observe that, for N sufficiently large, the function

$$f(x) = x^2 e^{-\frac{x^2}{2}},$$

is decreasing for $x > \delta\sqrt{N}$. Therefore, for N sufficiently large, we have

$$\begin{aligned}
\sum_{r=\lceil \delta\sigma_N\sqrt{N}\rceil}^{\infty} \frac{r^2}{\sigma_N^2} e^{-\frac{r^2}{2\sigma_N^2}} &\leq \int_{\lceil \delta\sigma_N\sqrt{N}\rceil - 1}^{\infty} \frac{x^2}{\sigma_N^2} e^{-\frac{x^2}{2\sigma_N^2}}\,\mathrm{d}x \\
&= \left[-xe^{-\frac{x^2}{2\sigma_N^2}}\right]_{\lceil \delta\sigma_N\sqrt{N}\rceil - 1}^{\infty} + \int_{\lceil \delta\sigma_N\sqrt{N}\rceil - 1}^{\infty} e^{-\frac{x^2}{2\sigma_N^2}}\,\mathrm{d}x \\
&= (\lceil \delta\sigma_N\sqrt{N}\rceil - 1)e^{-\frac{(\lceil \delta\sigma_N\sqrt{N}\rceil - 1)^2}{2\sigma_N^2}} \\
&\quad + \sigma_N\sqrt{2\pi}\int_{\lceil \delta\sigma_N\sqrt{N}\rceil - 1}^{\infty} \frac{1}{\sigma_N\sqrt{2\pi}} e^{-\frac{x^2}{2\sigma_N^2}}\,\mathrm{d}x
\end{aligned} \tag{34}$$

Recall the following standard Gaussian tail bound: if $X \sim \mathcal{N}(0, \sigma_N^2)$, we have

$$\mathbb{P}(X \geq t) \leq e^{-\frac{t^2}{2\sigma_N^2}}.$$

This is an elementary result; a simple proof is given in [28, Proposition 2.1.2]. Applying this Gaussian tail bound to quantity (34), we find that

$$\begin{aligned}
\sum_{r=\lceil \delta\sigma_N\sqrt{N}\rceil}^{\infty} \frac{r^2}{\sigma_N^2} e^{-\frac{r^2}{2\sigma_N^2}} &\leq (\lceil \delta\sigma_N\sqrt{N}\rceil - 1)e^{-\frac{(\lceil \delta\sigma_N\sqrt{N}\rceil - 1)^2}{2\sigma_N^2}} \\
&\quad + \sigma_N\sqrt{2\pi}e^{-\frac{(\lceil \delta\sigma_N\sqrt{N}\rceil - 1)^2}{2\sigma_N^2}} \\
&\leq \delta\sigma_N\sqrt{N}e^{-\frac{(\delta\sigma_N\sqrt{N}-1)^2}{2\sigma_N^2}} \qquad (35) \\
&\quad + \sigma_N\sqrt{2\pi}e^{-\frac{(\delta\sigma_N\sqrt{N}-1)^2}{2\sigma_N^2}}. \qquad (36)
\end{aligned}$$

Putting this all together, we obtain the following bound on quantity (33):

$$\begin{aligned}
\frac{1}{\sigma_N^2}\mathbb{E}\left((a_{0,N}^{(1)})^2 \mathbf{1}_{|a_{0,N}^{(1)}| \geq \delta\sigma_N\sqrt{N}}\right) &\leq \frac{2}{\sigma_N\sqrt{2\pi} + O(1)}\sigma_N(\delta\sqrt{N} + \sqrt{2\pi})e^{-\frac{(\delta\sigma_N\sqrt{N}-1)^2}{2\sigma_N^2}} \\
&= \frac{2}{\sqrt{2\pi} + O(1/\sigma_N)}(\delta\sqrt{N} + \sqrt{2\pi})e^{-\frac{(\delta\sigma_N\sqrt{N}-1)^2}{2\sigma_N^2}}.
\end{aligned}$$

Since $\delta > 0$, the exponential term dominates and this converges to 0 as $N \to \infty$.

Since Lindeberg's conditions hold for the triangular array $\{\xi_{N,r}/\gamma_N\}$ we conclude that

$$\frac{a}{\gamma_N\sigma_N\sqrt{N}}\Re(\lambda_{N,s}^{(1)}) + \frac{b}{\gamma_N\sigma_N\sqrt{N}}\Im(\lambda_{N,s}^{(1)}) \xrightarrow{d} \mathcal{N}(0,1),$$

for any $a, b \in \mathbb{R}$ with $a^2 + b^2 = 1$. Since $\gamma_N \to 1$ as $N \to \infty$, it follows that

$$\frac{a}{\sigma_N\sqrt{N}}\Re(\lambda_{N,s}^{(1)}) + \frac{b}{\sigma_N\sqrt{N}}\Im(\lambda_{N,s}^{(1)}) \xrightarrow{d} \mathcal{N}(0,1),$$

for any $a, b \in \mathbb{R}$ with $a^2 + b^2 = 1$. By the Cramér-Wold device,

$$\left(\frac{1}{\sigma_N\sqrt{N}}\Re(\lambda_{N,s}^{(1)}), \frac{1}{\sigma_N\sqrt{N}}\Im(\lambda_{N,s}^{(1)})\right) \xrightarrow{d} (X, Y),$$

where X, Y are independent random variables with distribution $\mathcal{N}(0,1)$. By the continuous mapping theorem, for any given s, we have, as $N \to \infty$,

$$\frac{1}{N\sigma_N^2}|\lambda_{N,s}^{(1)}|^2 \xrightarrow{d} \chi_2^2. \qquad \square$$

We can finally obtain a (probable) lower bound on the radius of correct decryption. Recall that we have correct decryption, provided that

$$||\varepsilon'||_\infty < R_{Dec} = \frac{\min_s \sqrt{|\lambda_{N,s}^{(1)}|^2 + |\lambda_{N,s}^{(2)}|^2}}{4N}.$$

Proposition 8. *Assume that* $\frac{1}{N\sigma_N^2}|\lambda_{N,s}^{(1)}|^2, \frac{1}{N\sigma_N^2}|\lambda_{N,s}^{(2)}|^2 \sim \chi_2^2$. *Then*

$$\mathbb{P}\left(R_{Dec} \geq \frac{\sigma_N}{4N^{3/4}}\right) \geq \frac{3}{4}.$$

Proof. First note that $R_{Dec} \geq \frac{\sigma_N}{4N^{3/4}}$ if, and only if,

$$\sqrt{|\lambda_{N,s}^{(1)}|^2 + |\lambda_{N,s}^{(2)}|^2} \geq \sigma_N \sqrt[4]{N},$$

for all s. Therefore,

$$\mathbb{P}\left(R_{Dec} < \frac{\sigma_N}{4N^{3/4}}\right) = \mathbb{P}\left(|\lambda_{N,s}^{(1)}|^2 + |\lambda_{N,s}^{(2)}|^2 < \sigma_N^2\sqrt{N} \text{ for some } s\right).$$

Applying the union bound, we obtain

$$\mathbb{P}\left(R_{Dec} < \frac{\sigma_N}{4N^{3/4}}\right) \leq N\mathbb{P}\left(|\lambda_{N,0}^{(1)}|^2 + |\lambda_{N,0}^{(2)}|^2 < \sigma_N^2\sqrt{N}\right). \tag{37}$$

Since $\lambda_{N,s}^{(1)}$ and $\lambda_{N,s}^{(2)}$ are independent, the right hand side of inequality (37) is upper bounded by

$$N\mathbb{P}\left(|\lambda_{N,0}^{(1)}|^2 < \sigma_N^2\sqrt{N}\right)^2 = N\mathbb{P}\left(\frac{1}{N\sigma_N^2}|\lambda_{N,0}^{(1)}|^2 < \frac{1}{\sqrt{N}}\right)^2. \tag{38}$$

Let $f(x)$ denote the density of the χ_2^2 distribution at x for $x \in \mathbb{R}_{\geq 0}$. Note that $f(x)$ is decreasing on $\mathbb{R}_{\geq 0}$ and that $f(0) = \frac{1}{2}$. Combining this fact with inequalities (37) and (38), we find that

$$\mathbb{P}\left(R_{Dec} < \frac{\sigma_N}{4N^{3/4}}\right) \leq N\left(\frac{1}{2\sqrt{N}}\right)^2 = \frac{1}{4}.$$

In other words,

$$\mathbb{P}\left(R_{Dec} \geq \frac{\sigma_N}{4N^{3/4}}\right) \geq \frac{3}{4}. \qquad \square$$

Combining (7) and (8), we obtain the following corollary:

Corollary 1. Let $\sigma_N \geq 4N^{3/4}$ and sample $a_r^{(1)}, a_r^{(2)} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma_N}$ for $0 \leq r < N$. Form the element

$$\alpha = \sum_{r=0}^{N-1} a_r^{(1)} \zeta_n^r + j \sum_{r=0}^{N-1} a_r^{(2)} \zeta_n^r.$$

Let $\{\lambda_N^r : 0 \leq r < 2N\}$ be the eigenvalues of M_α and define the random variable

$$R_{Dec} = \frac{\min_r |\lambda_N^r|}{4N}.$$

Then

$$\mathbb{P}\left(R_{Dec} \geq \frac{\sigma_N}{4N^{3/4}}\right) \gtrsim \frac{3}{4}.$$