



FuLeakage: Breaking FuLeeca by Learning Attacks

Felicitas Hörmann^{1,2}  and Wessel van Woerden³ 

¹ Institute of Communications and Navigation, German Aerospace Center (DLR),
Oberpfaffenhofen–Wessling, Germany

`felicitas.hoermann@dlr.de`

² School of Computer Science, University of St. Gallen, St. Gallen, Switzerland

³ Univ. Bordeaux, CNRS, Inria, Bordeaux INP, IMB, Talence, France

`wessel.van-woerden@math.u-bordeaux.fr`

Abstract. FuLeeca is a signature scheme submitted to the recent NIST call for additional signatures. It is an efficient hash-and-sign scheme based on quasi-cyclic codes in the Lee metric and resembles the lattice-based signature FALCON. FuLeeca proposes a so-called concentration step within the signing procedure to avoid leakage of secret-key information from the signatures. However, FuLeeca is still vulnerable to learning attacks, which were first observed for lattice-based schemes. We present three full key-recovery attacks by exploiting the proximity of the *code*-based FuLeeca scheme to *lattice*-based primitives.

More precisely, we use a few signatures to extract an $n/2$ -dimensional circulant sublattice from the given length- n code, that still contains the exceptionally short secret-key vector. This significantly reduces the classical attack cost and, in addition, leads to a full key recovery in quantum-polynomial time. Furthermore, we exploit a bias in the concentration procedure to classically recover the full key for any security level with at most 175,000 signatures in less than an hour.

1 Introduction

Most of today’s asymmetric cryptosystems are based on discrete-logarithm problems or integer factorization and will not withstand powerful quantum computers. Thus, the U.S. National Institute of Standards and Technology (NIST) currently aims to standardize cryptographic schemes that resist conventional as well as quantum attacks. The project started with the first call for post-quantum algorithms in 2016 and now arrived at the fourth and last round [25]. Since many of the remaining signature schemes share similar underlying security primitives, NIST opened up a call for additional signatures in 2023 [26] to foster diversity.

One of the new proposals is FuLeeca [32,33], which is the first scheme relying on coding-theoretical problems in the Lee metric. FuLeeca adopts a hash-and-sign approach and thus generates signatures as codewords that are close to the hashed message. Then, *any* basis of the code allows the recipient to verify the signature, whereas the signer needs access to the secret key, i.e., a *good* basis for decoding the signature from the target hash.

The hash-and-sign paradigm is the basis for code-based schemes such as CFS [11] and Wave [5], but it is also well-known in lattice-based cryptography. Unfortunately, the first lattice proposals adopting this strategy, namely GGH [19] and NTRUSIGN [20], were fully broken by learning attacks [18, 27]. In particular, already a few hundred NTRUSIGN signatures leak enough information to recover the secret key, as their generation heavily depends on the secret lattice basis. Several heuristic approaches tried to prevent leakage [21, 22, 28, 35] but failed in the end [14, 23, 37]. In contrast, signatures generated with the GPV framework [17] *provably* do not leak any information about the secret key and thus resist learning-type attacks. Instantiations of the GPV framework are e.g. the code-based scheme Wave [5] and the lattice-based signature FALCON [31], and NIST selected the latter for standardization [25].

FuLeeca does *not* adopt the GPV framework and provides *no* proof of non-leakage. Instead, a heuristic *concentration step* is added to the signing algorithm to prevent leakage. Unfortunately, this countermeasure is not enough.

Table 1: Overview of the presented attacks against FuLeeca.

	few signatures ($\ll 100$)	many signatures ($\leq 175,000$)
classical attack	leaked-sublattice attack (reduced security) → section 4	learning attack (full break) → section 6
quantum attack	ideal-structure attack (full break) → section 5	← see this attack

Contributions. We present three types of attacks against the signature scheme FuLeeca, which all profit from the proximity of FuLeeca to known lattice-based constructions. This showcases once more the close connection between code- and lattice-based cryptography. Table 1 classifies the three attacks and provides pointers to the corresponding sections. In the following, we give a short summary of each approach:

Firstly, we present a *leaked-sublattice attack*. The idea of applying basis reduction to the construction-A lattice obtained from the public key was already considered in the FuLeeca specification but did not influence the parameter choices as it is less efficient than other combinatorial attacks. We observe however that all signatures are part of a lower-dimensional sublattice generated by the short secret vectors. Moving to this lower-dimensional sublattice, for which a basis can be computed from a small sample of signatures, gives a drastic speedup. The attack reduces the security level of the FuLeeca-I, FuLeeca-III, and FuLeeca-V parameter sets from 160, 224, and 288 bits to only 111, 155, and 199 bits, respectively.

Secondly, we perform an *ideal-structure attack* by realizing that the first half of the secret vector is an unusually short element in a circulant lattice. Again,

a basis for this circulant lattice can be computed from a few signatures. This enables a polynomial-time quantum attack in the FuLeeca setting, thus giving a full quantum key recovery for all three security levels.

Thirdly, we set up a *learning attack* which yields a full key recovery when enough signature vectors are known. The attack exploits a bias in the signing procedure, which lets every signature leak information about the secret key. In practice, we recover the full secret key of FuLeeca-I, FuLeeca-III, and FuLeeca-V instances with only 90,000, 175,000, and 175,000 signatures, respectively. The observed success rate is shown in Figure 1 and is indeed higher for the supposedly more secure FuLeeca-V instances than for the FuLeeca-III samples.




Fig. 1: Success rate of the learning attack with respect to the number of available signatures. Averaged over 50 instantiations for each parameter set.

In summary, we can perform a full key recovery in polynomial time if a reasonable amount of signatures was collected, and in quantum-polynomial time for a small sample of signatures. This is a full break of FuLeeca for all proposed parameter sets, regardless of the availability of quantum computers.

2 Preliminaries

We aim at making this paper accessible for readers from both the lattice community and the coding community. Thus, we give rather detailed preliminaries to provide a mostly self-contained presentation of our results.

Generic notation. Let p be an odd prime and denote the finite field of order p by $\mathbb{F}_p \cong \mathbb{Z}/p\mathbb{Z}$. Note that FuLeeca, and thus this work, fixes $p = 65,521$ in all parameter sets. Vectors are considered as *row* vectors and are denoted in lower-case bold, as e.g. \mathbf{a} , while matrices are denoted in upper-case bold, as e.g. \mathbf{A} . We identify \mathbb{F}_p with $\{-\frac{p-1}{2}, \dots, \frac{p-1}{2}\}$ and implicitly lift vectors and matrices over \mathbb{F}_p to the corresponding objects over \mathbb{Z} .

For a real number $x \in \mathbb{R}$, we write $\lfloor x \rfloor$ for the largest integer $n \leq x$, $\lceil x \rceil$ for the smallest integer $n \geq x$, and $\lfloor x \rfloor \in \mathbb{Z}$ for the unique integer such that $x - \lfloor x \rfloor \in [-\frac{1}{2}, \frac{1}{2})$ holds. We use $\text{trunc}(x)$ to denote the rounding of x towards 0, i.e., $\text{trunc}(x) = \lfloor x \rfloor$ if $x \geq 0$, and $\text{trunc}(x) = \lceil x \rceil$ otherwise. The *sign* of x is denoted by $\text{sgn}(x) \in \{-1, 0, 1\}$, and defined as $\text{sgn}(x) = 1$ if $x > 0$, $\text{sgn}(x) = -1$ if $x < 0$, and $\text{sgn}(0) = 0$. For a vector \mathbf{a} or a matrix \mathbf{A} with coefficients in a set S , we write $f(\mathbf{a})$ or $f(\mathbf{A})$ to apply a function $f : S \rightarrow S$ coefficient-wise.

Circulant matrices. Let \mathbb{F} be any field. We define the *circular shift* of a vector $\mathbf{a} = (a_1, \dots, a_k) \in \mathbb{F}^k$ as $\text{shift}(\mathbf{a}) := (a_k, a_1, \dots, a_{k-1}) \in \mathbb{F}^k$. For $i \in \mathbb{N}$, the i -fold shift of \mathbf{a} is obtained by applying the shift-operator i times and denoted by $\text{shift}^i(\mathbf{a})$. We further construct a matrix $\text{Shift}(\mathbf{a}) \in \mathbb{F}^{k \times k}$ by using all possible shifts of a given vector $\mathbf{a} \in \mathbb{F}^k$ as its rows. Namely,

$$\text{Shift}(\mathbf{a}) := \begin{pmatrix} \mathbf{a} \\ \text{shift}(\mathbf{a}) \\ \vdots \\ \text{shift}^{k-1}(\mathbf{a}) \end{pmatrix} = \begin{pmatrix} a_1 & a_2 & \dots & a_k \\ a_k & a_1 & \dots & a_{k-1} \\ \vdots & \vdots & \ddots & \vdots \\ a_2 & a_3 & \dots & a_1 \end{pmatrix}.$$

The matrix $\text{Shift}(\mathbf{a})$ is called *circulant* and the identity $\text{shift}(\mathbf{a}) \cdot \text{Shift}(\mathbf{b}) = \text{shift}(\mathbf{b}) \cdot \text{Shift}(\mathbf{a})$ holds for any $\mathbf{a}, \mathbf{b} \in \mathbb{F}^k$.

2.1 Codes

A *linear code* \mathcal{C} of length n and dimension k is a k -dimensional subspace of \mathbb{F}_p^n . We usually represent it as the row space of a matrix and call any full-rank matrix $\mathbf{G} \in \mathbb{F}_p^{k \times n}$ with $\mathcal{C} = \langle \mathbf{G} \rangle_{\mathbb{F}_p}$ a *generator matrix* of \mathcal{C} . We can apply Gaussian elimination to transform a generator matrix, up to column permutations, into its systematic form $(\mathbf{I}_k \mid \mathbf{T})$, where $\mathbf{T} \in \mathbb{F}_p^{k \times (n-k)}$ and $\mathbf{I}_k \in \mathbb{F}_p^{k \times k}$ is the identity matrix. The classical weight that is considered for linear codes is the *Hamming weight*. It counts the number of nonzero entries of a vector, i.e.,

$$\text{wt}_H(\mathbf{x}) := |\{i \in \{1, \dots, n\} : x_i \neq 0\}| \quad \text{for any } \mathbf{x} = (x_1, \dots, x_n) \in \mathbb{F}_p^n.$$

Quasi-cyclic codes. A k -dimensional *quasi-cyclic code* of length $n = 2k$ has a generator matrix $\mathbf{G} = (\mathbf{A} \mid \mathbf{B}) \in \mathbb{F}_p^{k \times n}$ consisting of two circulant blocks $\mathbf{A} = \text{Shift}(\mathbf{a}) \in \mathbb{F}_p^{k \times k}$ and $\mathbf{B} = \text{Shift}(\mathbf{b}) \in \mathbb{F}_p^{k \times k}$. If \mathbf{A} is invertible over \mathbb{F}_p , the matrix $\mathbf{T} = \mathbf{A}^{-1}\mathbf{B}$ in the systematic form $(\mathbf{I}_k \mid \mathbf{T})$ of \mathbf{G} is also circulant.

Lee metric. The *Lee weight* of an element $x \in \mathbb{F}_p$ is defined as $\text{wt}_L(x) := |x|$, where we use the symmetric field representation $\mathbb{F}_p = \{-\frac{p-1}{2}, \dots, \frac{p-1}{2}\}$. The Lee weight can be extended to vectors additively, i.e.,

$$\text{wt}_L(\mathbf{x}) = \sum_{i=1}^n \text{wt}_L(x_i) \quad \text{for any } \mathbf{x} = (x_1, \dots, x_n) \in \mathbb{F}_p^n.$$

The metric induced by this weight, that is, $d_L(\mathbf{x}, \mathbf{y}) := \text{wt}_L(\mathbf{x} - \mathbf{y})$ for $\mathbf{x}, \mathbf{y} \in \mathbb{F}_p^n$, is called the *Lee metric* on \mathbb{F}_p^n .

Typical Lee vectors. Suppose one wants to sample a vector $\mathbf{x} \in \mathbb{F}_p^n$ uniformly at random from the Lee sphere

$$\mathcal{S}_L(t, n) := \{\mathbf{v} \in \mathbb{F}_p^n : \text{wt}_L(\mathbf{v}) = t\}$$

of vectors having a fixed Lee weight t . This is not as straightforward as e.g. for the Hamming metric, where sampling t uniform nonzero values and a random permutation suffices. A vector of Lee weight t can be obtained by first choosing a weight partition $t = t_1 + \dots + t_n$ with integers $0 \leq t_i \leq \frac{p-1}{2}$ and then returning a signed permutation of the vector $\mathbf{t} = (t_1, \dots, t_n)$. Uniform sampling over $\mathcal{S}_L(t, n)$ is possible by properly randomizing the chosen partition and the signed permutation [6].

However, randomizing the partition correctly is costly and therefore FuLeeca fixes a precomputed *typical partition* or *typical Lee vector* \mathbf{t} for the used length n , modulus p , and Lee weight t , and only randomizes by applying a signed permutation to it. In other words, FuLeeca samples uniformly from the set

$$T(\mathbf{t}) := \{\pi(\mathbf{t}) : \pi \text{ is a signed permutation}\}.$$

2.2 Lattices

A *lattice* $\mathcal{L} \subset \mathbb{R}^n$ is a discrete subgroup of the vector space \mathbb{R}^n . A prominent example is the integer lattice $\mathbb{Z}^n \subset \mathbb{R}^n$. More generally, let $\mathbf{B} \in \mathbb{R}^{k \times n}$ be the matrix whose rows are the \mathbb{R} -linearly independent vectors $\mathbf{b}_1, \dots, \mathbf{b}_k \in \mathbb{R}^n$. Then we define the lattice $\mathcal{L}(\mathbf{B})$ with *basis* \mathbf{B} by

$$\mathcal{L}(\mathbf{B}) := \mathbb{Z}^k \cdot \mathbf{B} := \left\{ \sum_{i=1}^k x_i \mathbf{b}_i : x_i \in \mathbb{Z} \quad \forall i \right\}.$$

We call k the *rank* or *dimension* of the lattice, and $\mathcal{L}(\mathbf{B})$ has full rank if $k = n$.

For lattices $\mathcal{L} \subset \mathbb{R}^n$, we usually consider the standard *Euclidean metric* in \mathbb{R}^n . The Euclidean inner product $\langle \mathbf{x}, \mathbf{y} \rangle$ for two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ is given by $\langle \mathbf{x}, \mathbf{y} \rangle := \sum_{i=1}^n x_i y_i$. It induces the Euclidean (or ℓ_2) norm by

$$\|\mathbf{x}\|_2 := \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle} \quad \text{for any } \mathbf{x} \in \mathbb{R}^n.$$

Alternatively, and related to the Lee metric, one can consider the metric induced by the ℓ_1 norm $\|\mathbf{x}\|_1 := \sum_{i=1}^n |x_i|$.

A lattice has a *volume* $\text{vol}(\mathcal{L}) := \sqrt{\det(\mathbf{B}\mathbf{B}^\top)}$, where \mathbf{B} is an arbitrary basis of \mathcal{L} . Note that the volume is independent of the considered basis despite the given definition because every basis $\tilde{\mathbf{B}}$ has the form $\tilde{\mathbf{B}} = \mathbf{U}\mathbf{B}$ for a unimodular matrix $\mathbf{U} \in \text{GL}_k(\mathbb{Z})$ with determinant $\det(\mathbf{U}) = \pm 1$.

The length of any shortest nonzero lattice vector is denoted by

$$\lambda_1(\mathcal{L}) := \min_{\mathbf{v} \in \mathcal{L}, \mathbf{v} \neq \mathbf{0}} \|\mathbf{v}\|_2$$

and called the *first minimum* of \mathcal{L} . Note that $\lambda_1(\mathcal{L})$ is a *positive* real because of the discrete nature of lattices. For a random lattice of rank k ,⁴ the first minimum $\lambda_1(\mathcal{L})$ is concentrated around $\text{gh}(\mathcal{L}) := \text{gh}(k) \cdot \text{vol}(\mathcal{L})^{1/k}$ with high probability, where $\text{gh}(k) := \text{vol}(\mathcal{B}_1^k)^{-1/k} \approx \sqrt{k/(2\pi e)}$ is the *Gaussian heuristic* and \mathcal{B}_1^k denotes the unit ball of dimension k . The Gaussian heuristic can be seen as the lattice analog of the Gilbert–Varshamov (GV) bound for codes.

Computing short lattice vectors. The *shortest-vector problem (SVP)* takes a lattice basis \mathbf{B} as input and asks to compute a shortest lattice vector, i.e., a $\mathbf{v} \in \mathcal{L}(\mathbf{B})$ with $\|\mathbf{v}\|_2 = \lambda_1(\mathcal{L}(\mathbf{B}))$. Finding a shortest vector of a high-dimensional lattice is in general a hard problem. State-of-the-art heuristic algorithms for SVP take both exponential time $2^{0.292k+o(k)}$ and memory $2^{0.208k+o(k)}$ for a k -dimensional lattice. However, finding a *short-enough* vector might suffice to break cryptographic schemes. Lattice-reduction algorithms such as BKZ [34] find such short-enough vectors with a trade-off between the vector length and the computational cost. BKZ requires an oracle that solves SVP *exactly* in β -dimensional lattices for $\beta < k$ and runs in time $2^{0.292\beta+o(k)}$. We say that BKZ runs with *blocksize* β and a higher blocksize heuristically recovers shorter vectors.

Heuristic Claim 1 (BKZ approximation [2]) *For a k -dimensional lattice \mathcal{L} , BKZ with blocksize $2 \leq \beta \leq k$ heuristically recovers a lattice vector $\mathbf{v} \in \mathcal{L}$ of length*

$$\|\mathbf{v}\|_2 \leq \text{gh}(\beta)^{\frac{k-1}{\beta-1}} \cdot \text{vol}(\mathcal{L})^{1/k} \approx \left(\frac{\beta}{2\pi e}\right)^{\frac{k-1}{2(\beta-1)}} \cdot \text{vol}(\mathcal{L})^{1/k}$$

in time $2^{0.292\beta+o(k)}$.

The special cases of the BKZ algorithm are $\beta = 2$, for which BKZ runs in polynomial time but gives an exponentially large approximation radius, and $\beta = k$, for which it runs in exponential time but recovers a shortest lattice vector. If the given lattice is not random, BKZ might be faster. For example, if the lattice contains an unusually short vector, i.e., a vector shorter than what the Gaussian heuristic predicts, BKZ can recover it with a blocksize $\beta \ll k$.

Heuristic Claim 2 (BKZ unusual SVP [2]) *Let \mathcal{L} be a lattice of dimension k and let $\mathbf{v} \in \mathcal{L}$ be an unusually short vector with $\|\mathbf{v}\|_2 \ll \text{gh}(\mathcal{L})$. Then, BKZ with blocksize β heuristically recovers \mathbf{v} if*

$$\sqrt{\frac{\beta}{k}} \cdot \|\mathbf{v}\|_2 < \text{gh}(\beta)^{\frac{2\beta-k-1}{\beta-1}} \cdot \text{vol}(\mathcal{L})^{1/k}.$$

For example, if $\|\mathbf{v}\|_2 = \text{gh}(\mathcal{L})/\Theta(\sqrt{k})$, then BKZ recovers the unusually short vector with blocksize $\beta = \frac{k}{2} + o(k)$ and thus in time $2^{\frac{0.292k}{2}+o(k)}$ instead of $2^{0.292k+o(k)}$. The above estimate can be refined further, leading to accurate concrete predictions of the precise blocksize β that is needed to recover an unusually short lattice vector [4, 13, 29]. See [3] for a survey on these results.

⁴ The meaning of a *random* lattice is a bit more complex than for codes, but can be made rigorous by considering the unique normalized Haar measure on the space of lattices.

2.3 Ideal and log-unit lattices

Recall that an *ideal* \mathcal{I} of a ring R is an additive subgroup of R satisfying $R \cdot \mathcal{I} = \mathcal{I}$. Any ideal is already equipped with the additive properties of a lattice, but there is no a-priori notion of a distance. We now discuss two constructions of *ideal lattices*, their relation, and the log-unit lattice.

Coefficient embedding. Firstly, consider polynomial quotient rings of the form $R := \mathbb{Z}[x]/(h)$ for a monic polynomial $h \in \mathbb{Z}[x]$ of degree d . We define the *coefficient embedding* $\text{cf} : R \rightarrow \mathbb{Z}^d$, which maps any element $f(x) = \sum_{i=1}^d f_i x^{i-1} \in R$ to its coefficient vector $\text{cf}(f) := (f_1, \dots, f_d) \in \mathbb{Z}^d$. Since $\mathbb{Z}^d \subset \mathbb{R}^d$ is equipped with the Euclidean norm $\|\cdot\|_2$, this mapping induces a metric on R by setting $\|f\| := \|\text{cf}(f)\|_2$ for any $f \in R$. Thus, we call the image $\text{cf}(\mathcal{I}) \subset \mathbb{R}^d$ of any ideal $\mathcal{I} \subset R$ an ideal lattice. We will encounter an example of this construction in the following, where we consider the case $h(x) := x^k - 1$ leading to the quotient ring $R = \mathbb{Z}[x]/(x^k - 1)$. In this setting, the lattice $\text{cf}(\mathcal{I})$ constructed from any principal ideal $\mathcal{I} = (a)$ in R with generator $a = \sum_{i=0}^{k-1} a_i x^i \pmod{x^k - 1}$ has a circulant basis. This derives from the fact that $a, xa, \dots, x^{k-1}a$ is a \mathbb{Z} -basis of \mathcal{I} and multiplication with x^i in R coincides with the i -fold shift of the image vector in $\text{cf}(\mathcal{I})$. More precisely, $\text{cf}(x^i a) = \text{shift}^i(\text{cf}(a))$ holds for $i = 0, \dots, k-1$ and thus we obtain $\text{cf}(\mathcal{I}) = \mathcal{L}(\text{Shift}(\mathbf{a}))$ with $\mathbf{a} := \text{cf}(a)$.

Canonical embedding. Secondly, we choose R as the ring of integers \mathcal{O}_K of a number field $K := \mathbb{Q}[x]/(h)$ with $h \in \mathbb{Q}[x]$ being an irreducible monic polynomial of degree d . Note that K is naturally equipped with d field homomorphisms $\sigma_i : K \rightarrow \mathbb{C}$ with $i = 1, \dots, d$ that map x to distinct roots of h in \mathbb{C} . We call $\sigma : K \rightarrow \mathbb{C}^d$ with $f \mapsto (\sigma_1(f), \dots, \sigma_d(f))$ for any $f \in K$ the *canonical embedding* of K into $\mathbb{C}^d \simeq \mathbb{R}^{2d}$. We use it to transfer the Euclidean metric from \mathbb{R}^{2d} to K , similar as we did with the coefficient embedding cf in the first construction. Namely, we set $\|f\| := \|\sigma(f)\|_2$ for any $f \in K$, where $\sigma(f)$ on the right-hand side is implicitly understood as an element of $\mathbb{R}^{2d} \simeq \mathbb{C}^d$. Since σ preserves the additive and discrete structure of any nonzero ideal $\mathcal{I} \subset R$, we call the rank- d image $\sigma(\mathcal{I}) \subset \mathbb{C}^d$ an ideal lattice. One particular example of this construction, which we will come across later, is the case of cyclotomic fields. For any odd prime k , we consider the monic *cyclotomic polynomial*

$$\Psi_k(x) := (x^k - 1)/(x - 1) = 1 + x + \dots + x^{k-1} \in \mathbb{Z}[x]$$

which is irreducible over \mathbb{Q} and has degree $d = k - 1$. We obtain d -dimensional ideal lattices by applying the canonical embedding to nonzero ideals in the ring of integers $\mathcal{O}_k := \mathbb{Z}[x]/(\Psi_k)$ of the *cyclotomic number field* $K := \mathbb{Q}[x]/(\Psi_k)$.

Circulant and cyclotomic ideals. The two examples we considered are in fact closely related. Namely, for a prime k , the quotient ring $\mathbb{Z}[x]/(x^k - 1)$ splits into

$$\mathbb{Z}[x]/(x^k - 1) \cong \mathbb{Z}[x]/(\Psi_k) \times \mathbb{Z}[x]/(x - 1)$$

by means of the isomorphism $f \mapsto (f \pmod{\Psi_k}, f \pmod{x-1})$. As a result, we can map any circulant ideal $\mathcal{I} \subset \mathbb{Z}[x]/(x^k - 1)$ of rank k onto a cyclotomic ideal $\mathcal{I} \pmod{\Psi_k} \subset \mathbb{Z}[x]/(\Psi_k)$ of rank $k - 1$, and onto an ideal $\mathcal{I} \pmod{x-1} \subset \mathbb{Z}[x]/(x-1) \cong \mathbb{Z}$ of rank 1. An important property of the first part of the mapping is that the geometry of the ideal lattice $\text{cf}(\mathcal{I})$ via the coefficient embedding is, up to scaling, very close to the ideal lattice $\sigma(\mathcal{I} \pmod{\Psi_k})$ via the canonical embedding [8]. So we can easily move from our circulant lattice to an ideal lattice over a cyclotomic number ring. This will become useful later because some generally hard lattice problems can be solved in quantum-polynomial time for principal ideal lattices over cyclotomic number rings.

Log-unit lattice. We recall one last embedding that is useful for the multiplicative structure of units. Consider again a number ring \mathcal{O}_K of a number field K of degree d with canonical embedding σ given by $\sigma_1, \dots, \sigma_d$. We define the *logarithmic embedding* $\text{Log} : K^\times \rightarrow \mathbb{R}^d$ by

$$f \mapsto (\log(|\sigma_1(f)|), \dots, \log(|\sigma_d(f)|)) \in \mathbb{R}^d.$$

Note that multiplication in $K^\times = K \setminus \{0\}$ corresponds to addition after the logarithmic embedding, i.e.,

$$\text{Log}(u \cdot v) = \text{Log}(u) + \text{Log}(v) \quad \text{for all } u, v \in K^\times.$$

Now consider the subgroup of units \mathcal{O}_K^\times of \mathcal{O}_K . By Dirichlet's unit theorem the logarithmic embedding $\text{Log}(\mathcal{O}_K^\times)$ is a lattice of a certain rank, named the *log-unit lattice*. Going back to our example, the log-unit lattice $\text{Log}(\mathcal{O}_k^\times)$ of the cyclotomic number field $K := \mathbb{Q}[x]/(\Psi_k)$ has rank $(k - 3)/2$ for all odd primes k . Because $|N(u)| := \prod_{i=1}^d |\sigma_i(u)| = 1$ holds for all units $u \in \mathcal{O}_K^\times$, we obtain that the coefficients of $\text{Log}(u)$ sum up to 0. In other words, $\text{Log}(\mathcal{O}_K^\times)$ lies in the subspace orthogonal to $(1, \dots, 1)$. Furthermore, any root of unity $\xi \in K$ satisfies $|\sigma_i(\xi)| = 1$ for all $i = 1, \dots, d$ and thus $\text{Log}(\xi) = \mathbf{0}$, i.e., ξ is contained in the kernel of Log . In particular, Log maps any product $\xi \cdot u$ of a unit $u \in K^\times$ and a root of unity $\xi \in K^\times$ to $\text{Log}(u)$, which means that u can only be recovered from its image $\text{Log}(u)$ up to a root of unity.

3 FuLeeca and lattice-based cryptography

FuLeeca is a hash-and-sign signature scheme based on codes in the Lee metric. The hash-and-sign paradigm is a well-known technique to design digital signatures and has been used for constructions involving both lattices and codes. The main idea is to interpret the hash of the given message as a target point in an ambient space in which a certain selected code or lattice is defined. As the signer knows a suitable *good* basis of the code or the lattice, they can easily find a codeword or a lattice point that is close to the target with respect to the corresponding metric and use it as the signature. Note that even a publicly known

bad basis, which fully hides the structure of the chosen lattice or code, allows the verifier to efficiently validate the signature by checking if the signature is a lattice point or a codeword and if it is close to the hash of the signed message. The obstacle for an adversary trying to forge a signature is however the hardness of finding a close codeword or lattice point without having access to a good basis. These problems are known as the syndrome-decoding problem (SDP) and the closest-vector problem (CVP), respectively.

Hash-and-sign-based examples from code-based cryptography start with the CFS scheme [11] and include the current NIST proposal Wave [5]. In terms of lattices, the most prominent examples of hash-and-sign signatures are GGH [19], NTRUSIGN [20], and FALCON [31].

Since the Lee metric is closely related to the Euclidean $\|\cdot\|_2$ -norm, not only the overall hash-and-sign structure of FuLeeca coincides with known lattice proposals, but also the underlying notion of distance is similar. We use the remainder of this section to recall the design principles of FuLeeca and describe how it can be interpreted in the context of lattice-based cryptography. This allows us to identify similarities and differences with respect to GGH, NTRUSIGN, and FALCON. Our observations will enable us to apply lattice techniques to attack FuLeeca later in this paper.

3.1 FuLeeca

The NIST submission of the FuLeeca signature scheme [32] is based on the publication [33]. It comes with three sets of parameters that can be found in Table 2. As the chosen parameter sets correspond to the NIST security levels I, III, and V, we refer to them as FuLeeca-I, FuLeeca-III, and FuLeeca-V, respectively.

Key generation. Both FuLeeca keys are generator matrices of the same quasi-cyclic code $\mathcal{C} \subset \mathbb{F}_p^n$ of dimension k and length $n = 2k$. The secret key $\mathbf{G}_{\text{sec}} = (\mathbf{A} \mid \mathbf{B}) \in \mathbb{F}_p^{k \times n}$ consists of two circulant blocks $\mathbf{A} \in \text{GL}(k, \mathbb{F}_p)$ and $\mathbf{B} \in \mathbb{F}_p^{k \times k}$, and thus captures the quasi-cyclic structure of \mathcal{C} . In contrast, the public key $\mathbf{G}_{\text{pub}} = (\mathbf{I}_k \mid \mathbf{A}^{-1}\mathbf{B})$ is the systematic form of \mathbf{G}_{sec} and hides the secret key.

As depicted in Algorithm 1, the secret key \mathbf{G}_{sec} is chosen by picking two random signed permutations \mathbf{a} and \mathbf{b} of a typical Lee vector $\mathbf{t} \in \mathbb{F}_p^k$ of Lee weight w_{key} and setting $\mathbf{A} := \text{Shift}(\mathbf{a})$ and $\mathbf{B} := \text{Shift}(\mathbf{b})$. The sampling of \mathbf{a} is repeated until \mathbf{A} is invertible to ensure that the systematic form of \mathbf{G}_{sec} , that is, \mathbf{G}_{pub} , can be computed.

Algorithm 1: FuLeeca key generation [32].

Input : FuLeeca parameter set,
typical Lee vector \mathbf{t} .
Output: Public key \mathbf{G}_{pub} ,
secret key \mathbf{G}_{sec} .

```

1 do
2    $\mathbf{a} \xleftarrow{\$} T(\mathbf{t})$ 
3    $\mathbf{A} = \text{Shift}(\mathbf{a})$ 
4 while  $\mathbf{A} \notin \text{GL}(k, \mathbb{F}_p)$ 
5  $\mathbf{b} \xleftarrow{\$} T(\mathbf{t})$ 
6  $\mathbf{B} = \text{Shift}(\mathbf{b})$ 
7  $\mathbf{G}_{\text{sec}} = (\mathbf{A} \mid \mathbf{B})$ 
8  $\mathbf{G}_{\text{pub}} = (\mathbf{I}_k \mid \mathbf{A}^{-1}\mathbf{B})$ 
9 return  $\mathbf{G}_{\text{pub}}, \mathbf{G}_{\text{sec}}$ 

```

Table 2: Proposed FuLeeca parameter sets [32].

Parameter Set	p	n	w_{key}	w_{sig}/n	s	n_{con}
FuLeeca-I	65,521	1,318	31,102	982.8	3/64	100
FuLeeca-III	65,521	1,982	46,552	982.8	9/256	90
FuLeeca-V	65,521	2,638	61,918	982.8	3/128	178

Signature generation. The signer has access to the secret key \mathbf{G}_{sec} and wishes to sign a message \mathbf{m} . The signing procedure of FuLeeca is illustrated in Algorithm 2 and starts with applying hash functions and adding randomness to \mathbf{m} , i.e., with the *hash* part of the hash-and-sign design paradigm. We write $\text{Hash}(\mathbf{m}, \text{salt})$ to denote the adopted two-stage hashing process, that outputs a vector $\mathbf{c} \in \{\pm 1\}^n$ and is described in detail in the FuLeeca specification [32].

Now the *sign* part takes place. It consists of two steps, as is highlighted by the usage of the functions `simpleSign` and `concentrate` in Algorithm 2. These are followed by final checks to decide whether the signature is accepted or the whole process is repeated.

Algorithm 2: FuLeeca signature generation [32].⁵

Input : FuLeeca parameter set, secret key $\mathbf{G}_{sec} \in \mathbb{F}_p^{k \times n}$, message \mathbf{m} .
Output: Signature $\mathbf{v} \in \mathbb{F}_p^n$, salt salt .

```

1 Repeat
2   salt  $\xleftarrow{\$}$   $\mathbb{F}_2^{256}$ 
3    $\mathbf{c} = \text{Hash}(\mathbf{m}, \text{salt}) \in \{\pm 1\}^n$ 
4    $\mathbf{v} = \text{simpleSign}(\mathbf{c}, \mathbf{G}_{sec})$ 
5    $\mathbf{v} = \text{concentrate}(\mathbf{c}, \mathbf{v}, \mathbf{G}_{sec})$ 
6   if  $w_{sig} - 2w_{key} < \text{wt}_L(\mathbf{v}) \leq w_{sig}$ 
7   and  $\text{LMP}(\mathbf{v}, \mathbf{c}) \geq \lambda + 64$  then
8     return salt,  $\mathbf{v}$ 
9   end
10 simpleSign( $\mathbf{c}, \mathbf{G}_{sec}$ ):
11   for  $i = 1, \dots, k$  do
12      $x_i = \text{trunc}(s \cdot \text{ipm}(\mathbf{g}_i, \mathbf{c}))$ 
13   end
14    $\mathbf{x} = (x_1, \dots, x_k)$ 
15    $\mathbf{v} = \mathbf{x}\mathbf{G}_{sec}$ 
16   return  $\mathbf{v}$ 
17 concentrate( $\mathbf{c}, \mathbf{v}, \mathbf{G}_{sec}$ ):
18    $\mathcal{A} = \{\pm 1, \dots, \pm k\}$ , lf = 1
19   for  $j = 1, \dots, n_{con}$  do
20      $\mathbf{v}' = \mathbf{0}$ 
21     for  $i = 1, -1, \dots, k, -k$  do
22        $\mathbf{v}'' = \mathbf{v} + \text{sgn}(i) \cdot \mathbf{g}_{|i|}$ 
23       if  $|\text{LMP}(\mathbf{v}'', \mathbf{c}) - (\lambda + 65)| \leq$ 
24          $|\text{LMP}(\mathbf{v}', \mathbf{c}) - (\lambda + 65)|$  and
25         ( $i \in \mathcal{A}$  or lf = 0) then
26          $\mathbf{v}' = \mathbf{v}'', i' = i$ 
27       end
28     end
29      $\mathbf{v} = \mathbf{v}', \mathcal{A} = \mathcal{A} \setminus \{-i'\}$ 
30   if  $\text{wt}_L(\mathbf{v}) > w_{sig} - w_{key}$  then
31     lf = 0
32   else
33     lf = 1
34   end
35   return  $\mathbf{v}$ 

```

⁵ There are significant differences between the specification and the reference implementation of FuLeeca. In particular, the order of the loop on line 21 within the concentration procedure is important for our attack but not properly defined in the specification. Whenever specification and reference implementation differ, we follow the implementation.

Define the number of *sign matches* of two vectors $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{F}_p^n$ and $\mathbf{y} = (y_1, \dots, y_n) \in \mathbb{F}_p^n$ as

$$\text{mt}(\mathbf{x}, \mathbf{y}) := |\{i \in \{1, \dots, n\} : \text{sgn}(x_i) = \text{sgn}(y_i) \text{ and } x_i, y_i \neq 0\}|.$$

Moreover, the elementwise star product of \mathbf{x} and \mathbf{y} is $\mathbf{x} \star \mathbf{y} := (x_1 y_1, \dots, x_n y_n) \in \mathbb{F}_p^n$ and we define

$$\text{ipm}(\mathbf{x}, \mathbf{y}) := \text{mt}(\mathbf{x}, \mathbf{y}) - \frac{\text{wt}_H(\mathbf{x} \star \mathbf{y})}{2}.$$

Note that $2 \cdot \text{ipm}(\mathbf{x}, \mathbf{y}) = \langle \text{sgn}(\mathbf{x}), \text{sgn}(\mathbf{y}) \rangle$ applies, where $\langle \cdot, \cdot \rangle$ denotes the Euclidean scalar product. Therefore, we call $\text{ipm}(\mathbf{x}, \mathbf{y})$ the scaled number of *inner-product matches* of \mathbf{x} and \mathbf{y} .

The first signing step is *simple signing* which maps the hash output $\mathbf{c} \in \{\pm 1\}^n$ to a suitable vector $\mathbf{x} \in \mathbb{F}_p^k$ depending on the inner-product matches between \mathbf{c} and the rows of \mathbf{G}_{sec} , i.e., $x_i = \text{trunc}(s \cdot \text{ipm}(\mathbf{g}_i, \mathbf{c}))$ for some parameter s and $i = 1, \dots, k$. Then, it generates a signature $\mathbf{v} = \mathbf{x} \mathbf{G}_{\text{sec}}$ by interpreting the obtained \mathbf{x} as coordinates with respect to the secret basis $\mathbf{g}_1, \dots, \mathbf{g}_k$ of the code $\mathcal{C} = \langle \mathbf{G}_{\text{sec}} \rangle_{\mathbb{F}_p}$. Interestingly, the simple-signing procedure can be expressed as

$$\mathbf{v} = \text{trunc}\left(\frac{s}{2} \cdot \mathbf{c} \cdot \text{sgn}\left(\mathbf{G}_{\text{sec}}^\top\right)\right) \cdot \mathbf{G}_{\text{sec}}.$$

The second signing step is *concentrating* and aims at shifting the Lee weight of the simple signature \mathbf{v} and the number of sign matches $\text{mt}(\mathbf{c}, \mathbf{v})$ to prescribed intervals. This is achieved by a trial-and-error-like process that successively adds or subtracts rows of the secret generator matrix \mathbf{G}_{sec} and then checks if the result was improved. Note that doing this for the i -th row \mathbf{g}_i corresponds to adapting the coefficient x_i to $x_i \pm 1 \pmod{p}$ for $i = 1, \dots, k$. The quality of the signature is measured by its Lee weight and the *logarithmic matching probability (LMP)*, which is defined for two fixed vectors $\mathbf{v} \in \mathbb{F}_p^n$ and $\mathbf{c} \in \{\pm 1\}^n$ as

$$\text{LMP}(\mathbf{v}, \mathbf{c}) := -\log_2(\mathbb{P}[\text{mt}(\mathbf{v}, \mathbf{y}) = \text{mt}(\mathbf{v}, \mathbf{c})]).$$

Here, $\mathbf{y} \in \{\pm 1\}^n$ is chosen uniformly at random and $\mathbb{P}[\text{mt}(\mathbf{v}, \mathbf{y}) = \text{mt}(\mathbf{v}, \mathbf{c})]$ denotes the probability that \mathbf{v} has as many sign matches with \mathbf{y} as with \mathbf{c} . Note that a large value of $\text{LMP}(\mathbf{v}, \mathbf{c})$ indicates that that number of sign matches $\text{mt}(\mathbf{v}, \mathbf{c})$ is unusual, i.e., that the number of sign matches is significantly lower or higher than its expectation $\text{wt}_H(\mathbf{v} \star \mathbf{c})/2$ for random \mathbf{c} . The behavior of $\text{ipm}(\mathbf{c}, \mathbf{v})$ is similar, and $\text{ipm}(\mathbf{c}, \mathbf{v})$ has a large absolute value in this case as well.

The concentration process is repeated n_{con} times, where n_{con} is a preset parameter. Afterward, the obtained signature is accepted if its Lee weight is within the desired interval $(w_{\text{sig}} - 2w_{\text{key}}, w_{\text{sig}}]$ and the LMP between \mathbf{v} and \mathbf{c} is at least the bit-security level λ plus a security margin of 64 bits.

Signature verification. The verification procedure checks that the signature \mathbf{v} is part of the code and has small Lee weight $\text{wt}_L(\mathbf{v}) \leq w_{\text{sig}}$. Moreover, the signs of \mathbf{v} have to match sufficiently with \mathbf{c} , i.e., satisfy $\text{LMP}(\mathbf{v}, \mathbf{c}) \geq \lambda + 64$.

3.2 Interpreting FuLeeca as a lattice-based scheme

We can port the coding-theoretic ideas used in FuLeeca to the lattice setting by applying *construction A* [10, p. 137] to the considered code $\mathcal{C} = \langle \mathbf{G}_{\text{sec}} \rangle_{\mathbb{F}_p} \subset \mathbb{F}_p^n$ of length n and dimension k . Namely, we consider the lattice

$$\mathcal{L}_1 := \mathcal{C} + p\mathbb{Z}^n = \{ \mathbf{v} \in \mathbb{Z}^n : \mathbf{v} \pmod{p} \in \mathcal{C} \} \subset \mathbb{R}^n. \quad (1)$$

Recall that we identify \mathbb{F}_p with $\{-\frac{p-1}{2}, \dots, \frac{p-1}{2}\}$ and that we implicitly lift the entries of vectors and matrices from \mathbb{F}_p to \mathbb{Z} . The public generator matrix $\mathbf{G}_{\text{pub}} = (\mathbf{I}_k \mid \mathbf{A}^{-1}\mathbf{B})$ of \mathcal{C} , whose right side $\mathbf{A}^{-1}\mathbf{B}$ is computed modulo p , gives a public basis \mathbf{B}_{pub} of \mathcal{L}_1 . Namely,

$$\mathbf{B}_{\text{pub}} = \begin{pmatrix} \mathbf{I}_k & \mathbf{A}^{-1}\mathbf{B} \\ \mathbf{0} & p\mathbf{I}_{n-k} \end{pmatrix}$$

and this shows that \mathcal{L}_1 has full rank n and volume $\text{vol}(\mathcal{L}_1) = \det(\mathbf{B}_{\text{pub}}) = p^{n-k}$. The normalized volume $\text{vol}(\mathcal{L}_1)^{1/n} = p^{1-r}$ of the lattice is thus determined by the modulus p and the code rate $r := \frac{k}{n}$, whereas the rank of \mathcal{L}_1 equals the length of the code \mathcal{C} , as $p\mathbb{Z}^n \subset \mathcal{L}_1$ applies by construction. Since we will make use of different lattices throughout the paper, Figure 2 contains an overview of them, visualizes their connections, and indicates how to obtain suitable bases.

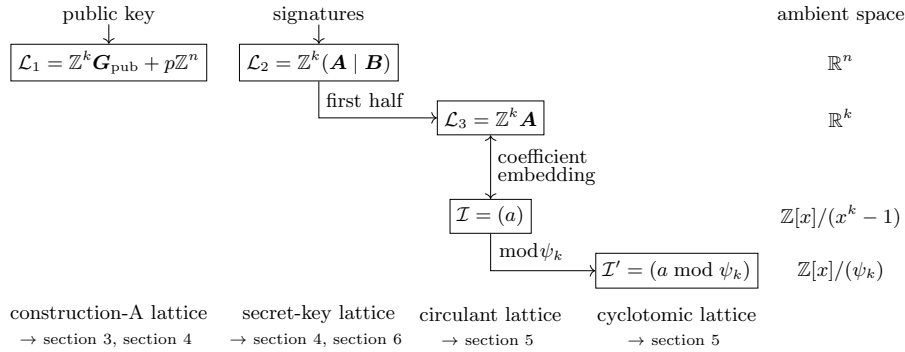


Fig. 2: Schematic overview of the lattices used in this paper and their connections.

Observe that the Lee metric on \mathcal{C} is closely related to the ℓ_1 metric on \mathcal{L}_1 . In particular, lifting a vector $\mathbf{x} \in \mathcal{C}$ with $\text{wt}_L(\mathbf{x}) = w$ implicitly to $\mathbf{x} \in \mathcal{L}_1$ yields $\|\mathbf{x}\|_1 = w$. The ℓ_1 norm is in fact close to the Euclidean ℓ_2 metric, to which most lattice algorithms are tailored. More precisely,

$$\|\mathbf{y}\|_2 \leq \|\mathbf{y}\|_1 \leq \sqrt{n} \cdot \|\mathbf{y}\|_2$$

holds for all $\mathbf{y} \in \mathbb{R}^n$ and we expect $\|\mathbf{y}\|_1 \sim \|\mathbf{y}\|_2$ for very sparse vectors and $\|\mathbf{y}\|_1 \sim \sqrt{n} \cdot \|\mathbf{y}\|_2$ for dense and balanced vectors. Short or close codewords of \mathcal{C} in the Lee metric are therefore directly related to short or close lattice vectors of \mathcal{L}_1 in the Euclidean metric.

In particular, we know that the short rows of the generator matrix \mathbf{G}_{sec} of \mathcal{C} correspond to k short vectors $\mathbf{g}_1, \dots, \mathbf{g}_k$ in \mathcal{L}_1 . As each \mathbf{g}_i has the form $(\text{shift}^i(\mathbf{a}) \mid \text{shift}^i(\mathbf{b}))$ and \mathbf{a} and \mathbf{b} are signed permutations of a precomputed typical vector \mathbf{t} , we know that $\|\mathbf{g}_i\|_2 = \|\mathbf{g}\|_2 = \sqrt{2} \cdot \|\mathbf{t}\|_2$ applies for $\mathbf{g} := (\mathbf{a} \mid \mathbf{b})$ and all $i = 1, \dots, k$. As can be seen in Table 3, this is close to the Gaussian heuristic of \mathcal{L}_1 , i.e., the expected first minimum of the lattice. The *good* generator matrix \mathbf{G}_{sec} thus gives a *good* basis of the sublattice $\mathcal{L}(\mathbf{G}_{\text{sec}}) \subset \mathcal{L}_1$.

FuLeeca uses this good basis to compute a signature \mathbf{v} that is short in the ℓ_1 norm and has a large absolute inner-product matching $|\text{ipm}(\mathbf{v}, \mathbf{c})| = \frac{1}{2} |\langle \text{sgn}(\mathbf{v}), \text{sgn}(\mathbf{c}) \rangle|$ with the hash $\mathbf{c} \in \{\pm 1\}^n$. Note that for a balanced vector $\mathbf{v} = (\pm v, \dots, \pm v)$ the inner-product matching correlates perfectly with the actual Euclidean inner product, i.e.,

$$\langle \mathbf{v}, \mathbf{c} \rangle = v \cdot \langle \text{sgn}(\mathbf{v}), \text{sgn}(\mathbf{c}) \rangle = 2v \cdot \text{ipm}(\mathbf{v}, \mathbf{c}).$$

More generally, we expect $\langle \mathbf{v}, \mathbf{c} \rangle \sim \frac{2\|\mathbf{v}\|_2}{\sqrt{n}} \cdot \text{ipm}(\mathbf{v}, \mathbf{c})$ for reasonably balanced vectors \mathbf{v} . The large absolute inner-product matching for a FuLeeca signature thus roughly corresponds to a large absolute Euclidean inner product between the signature \mathbf{v} and \mathbf{c} , i.e., \mathbf{v} and \mathbf{c} have a small angle. Furthermore, this is equivalent to saying that \mathbf{v} or $-\mathbf{v}$ is a close vector to $\frac{\|\mathbf{v}\|_2}{\|\mathbf{c}\|_2} \cdot \mathbf{c}$. So in terms of the Euclidean metric, a FuLeeca signature \mathbf{v} is valid when it is simultaneously small and close to the target obtained from the message.

Table 3: Lengths of the short secret-key vectors relative to the Gaussian heuristic.

Parameter Set	$\text{wt}_L(\mathbf{g})$	$\ \mathbf{g}\ _2$	$\ \mathbf{g}\ _2 / \text{gh}(\mathcal{L}_1)$	$\ \mathbf{g}\ _2 / \text{gh}(\mathcal{L}_2)$
FuLeeca-I	62,204	2,385.06	1.061	0.184
FuLeeca-III	93,104	2,976.61	1.079	0.150
FuLeeca-V	123,836	3,478.06	1.093	0.130

3.3 Comparing FuLeeca with known lattice-based schemes

We first compare FuLeeca to early lattice-based hash-and-sign signature schemes such as GGH [19] and NTRUSIGN [20]. GGH and NTRUSIGN both sign by hashing the message to a target in the space, and then using a good secret basis to compute a nearby lattice point. Recall from subsection 3.2 that a FuLeeca signature can also be interpreted as a lattice point close to some target hash. For FuLeeca however, the signature vector is also short, similar as in the recent lattice-based signature scheme HAWK [9, 15].

NTRUSIGN relies on structured NTRU lattices over the circulant quotient ring $\mathbb{Z}[x]/(x^k - 1)$, which correspond precisely to the lattice \mathcal{L}_1 in FuLeeca up to the choice of parameters. So both schemes use the same structure to achieve efficiency. Where FuLeeca uses a good basis \mathbf{G}_{sec} of the sublattice $\mathcal{L}(\mathbf{G}_{\text{sec}}) \subset \mathcal{L}_1$ to compute a nearby lattice point, NTRUSIGN and also FALCON [31] extend \mathbf{G}_{sec} to a full good basis \mathbf{B}_{sec} of the lattice \mathcal{L}_1 and use \mathbf{B}_{sec} for signature generation.

Unfortunately, both GGH and NTRUSIGN were broken by learning attacks, as each signature leaked some information about the secret key. For example, in the NTRUSIGN scheme the error \mathbf{e} between the target and the signature lattice vector is uniform over the parallelepiped $[-\frac{1}{2}, \frac{1}{2}]^n \cdot \mathbf{B}_{\text{sec}}$ of the secret basis \mathbf{B}_{sec} . Given enough signatures, it is possible to learn the parallelepiped, i.e., to recover the basis \mathbf{B}_{sec} [27]. For example, the computation $\mathbb{E}[\mathbf{e}^\top \mathbf{e}] = \frac{1}{12} \mathbf{B}_{\text{sec}} \mathbf{B}_{\text{sec}}^\top$ shows that the signatures leak the Gram matrix of the secret basis \mathbf{B}_{sec} . The full basis can be recovered by looking at higher moments.

Later variants of NTRUSIGN tried to mitigate learning attacks in various ways, for example by adding some extra noise to the decoding procedure. Unfortunately, these variants generally fell victim to similar learning attacks [14, 37], as the signatures still somehow depend directly on the secret information. Even recent hash-and-sign schemes, such as the submission PEREGRINE [35] in the Korean post-quantum-cryptography competition, are vulnerable to learning attacks [23]. In section 6, we show that FuLeeca has similar vulnerabilities.

An important breakthrough in hash-and-sign signature schemes was the GPV framework [17], that cleverly samples the nearby lattice point from a discrete Gaussian distribution around the target, using the secret basis. Because this distribution only depends on the lattice, they *proved* that the signatures do not leak any information about the secret basis, therefore mitigating any learning attack. FALCON [31], the lattice-based signature scheme that will be standardized by NIST [25], combines the hash-and-sign technique of NTRUSIGN and the GPV framework to prevent learning attacks. Similarly, the code-based hash-and-sign scheme Wave follows the GPV framework with ternary codes in the Hamming metric and uses appropriate rejection sampling to obtain a signature distribution that is independent of the secret code structure. In contrast, FuLeeca does not follow the GPV framework, and generally it remains an open question how to adapt the GPV framework to the Lee metric.

4 A leaked-sublattice attack

Recall that construction A allows to set up the lattice $\mathcal{L}_1 = \mathcal{C} + p\mathbb{Z}^n$ from the code \mathcal{C} spanned by the rows of the public generator matrix $\mathbf{G}_{\text{pub}} \in \mathbb{F}_p^{k \times n}$ as described in equation (1). As we explained in subsection 3.2, the rows $\mathbf{g}_1, \dots, \mathbf{g}_k$ of the secret generator matrix \mathbf{G}_{sec} are short in terms of the $\|\cdot\|_2$ -norm in the rank- n lattice \mathcal{L}_1 , i.e., they are between 1.06 and 1.1 times longer than the expected first minimum of \mathcal{L}_1 . Following Heuristic Claim 1, the BKZ lattice-reduction algorithm heuristically finds vectors of similar $\|\cdot\|_2$ -norm in time $2^{0.292\beta + o(n)}$ for $\beta \geq 0.95n$. Since the FuLeeca parameters are chosen such that the secret vectors $\mathbf{g}_1, \dots, \mathbf{g}_k$ are *not* unusually short in \mathcal{L}_1 , we cannot apply Heuristic Claim 2 to speed up the recovery.

The described approach is the only lattice-based attack that was considered in the FuLeeca specification and the security level was mainly determined by the cost of more efficient information-set-decoding (ISD) attacks [32, section 6]. In the following, we show that the transition to a suitable lower-dimensional

sublattice of \mathcal{L}_1 allows to achieve better attack complexities for lattice-reduction attacks. They outperform ISD attacks and lower the security for all parameter sets significantly.

4.1 Constructing a sublattice with unusually short vectors

The goal of this subsection is to describe a lower-dimensional sublattice \mathcal{L}_2 of \mathcal{L}_1 that still contains the secret vector $\mathbf{g} := (\mathbf{a} \mid \mathbf{b})$ and its quasi-circular shifts. This will lower the cost of recovering short vectors using BKZ. Moreover, we will see that the Euclidean norms of $\mathbf{g}_1, \dots, \mathbf{g}_k$ are unusually small for \mathcal{L}_2 , which makes the vectors even easier to recover.

Recall that each FuLeeca signature is given by a vector $\mathbf{v} = \mathbf{x}\mathbf{G}_{\text{sec}} \pmod{p}$ of low Lee weight. However, we observe experimentally that during the signature generation the coefficients stay small enough to *not* be reduced modulo p . This means that the equality $\mathbf{v} = \mathbf{x}\mathbf{G}_{\text{sec}}$ actually also holds over \mathbb{Z} , that is, if we interpret $\mathbf{v}, \mathbf{x} \in \left\{-\frac{p-1}{2}, \dots, \frac{p-1}{2}\right\}^n \subset \mathbb{Z}^n$ as integer vectors. Thus, all signature vectors belong to the lattice

$$\mathcal{L}_2 := \left\{ \sum_{i=1}^k z_i \mathbf{g}_i : z_i \in \mathbb{Z} \quad \forall i \right\} \subset \mathbb{R}^n \quad (2)$$

that is spanned by the rows of \mathbf{G}_{sec} .

We generated 5 million signatures per parameter set, which were split evenly over 50 FuLeeca key pairs in every case. All signature vectors belonged to \mathcal{L}_2 , which experimentally confirms the above observation. Furthermore, the largest observed coefficient size was 7,356, 8,564, and 7,616 for FuLeeca-I, FuLeeca-III, and FuLeeca-V instances, respectively. This is significantly smaller than $\frac{p-1}{2} = 32,760$. The described behavior of non-wrapping coefficients is explained by the fact that the simple-signing procedure creates an initial signature vector that is a small combination of the vectors $\mathbf{g}_1, \dots, \mathbf{g}_k$, which have small coefficients. As a result, no reduction modulo p takes place in simple signing. The concentration phase that follows adds vectors $\pm \mathbf{g}_i$ for $i = 1, \dots, k$ to the simple signature but furthermore tries to keep its Lee weight and thus its coefficients small, again leading to not wrapping the coefficients modulo p .

It is initially not clear that an attacker has access to a basis of \mathcal{L}_2 because \mathbf{G}_{sec} is secret and the knowledge of \mathbf{G}_{pub} only allows to recover a basis of \mathcal{L}_1 as described in subsection 3.2. However, a small sample of FuLeeca signatures is enough to recover a basis of \mathcal{L}_2 , which will be explained in subsection 4.2. Let us for now assume that we know a basis of \mathcal{L}_2 and describe how to recover the secret vector \mathbf{g} or one of its quasi-circular shifts under this assumption.

The lattice \mathcal{L}_2 has only rank $k = n/2$ and still contains the short vectors $\mathbf{g}_1, \dots, \mathbf{g}_k$. We can thus focus on the rank- k lattice \mathcal{L}_2 for the recovery of a secret vector. Compared to the rank- n lattice \mathcal{L}_1 , finding a shortest vector in \mathcal{L}_2 gives a reduced complexity of $2^{0.292k+o(k)}$.

Additionally, the vector \mathbf{g} is in fact an unusually short vector in \mathcal{L}_2 . Namely, \mathbf{g} is about a factor $\Theta(\sqrt{n})$ shorter than any vector one would expect in a lattice of

this rank and volume. This can be seen from the ratio of $\|\mathbf{g}\|_2$ and the Gaussian heuristic $\text{gh}(\mathcal{L}_2)$ of \mathcal{L}_2 in Table 3. As a result, a quasi-circular shift of \mathbf{g} can be recovered by the BKZ algorithm with $\beta = \frac{k}{2} + o(k) = \frac{n}{4} + o(n)$, leading to a time complexity of about $2^{\frac{0.292n}{4} + o(n)}$. In other words, we obtain a quartic speedup over the lattice attack considered in the FuLeeca specification [32]. A more precise value of β can be computed with concrete estimation scripts,⁶ leading to $\beta = 291$, $\beta = 448$, and $\beta = 603$ for the FuLeeca-I, FuLeeca-III, and FuLeeca-V parameters, respectively. This significantly reduces the claimed security levels of the selected parameter sets as shown in Table 4.

Table 4: Estimated security levels of the proposed FuLeeca parameter sets, claims according to [32, Table 1]. The new cost is based on the new blocksize estimate combined with the state-of-the-art cost model of [24] as implemented in the lattice estimator of [4].

Parameter Set	Claimed Security Level (in bits)	New Security Level (in bits)	Blocksize (β)
FuLeeca-I	160	111	291
FuLeeca-III	224	155	448
FuLeeca-V	288	199	603

4.2 Extracting the sublattice from FuLeeca signatures

The previous subsection showed that it is substantially cheaper to attack the sublattice \mathcal{L}_2 instead of \mathcal{L}_1 . But as the attacker does not have access to \mathbf{G}_{sec} , it is not trivial to get a basis of \mathcal{L}_2 . We now explain in more detail how to recover the lattice \mathcal{L}_2 from a few signatures.

Assuming signatures are not reduced modulo p , any signature vector $\mathbf{v} = (\mathbf{w} \mid \mathbf{y})$ defines a generating matrix $\mathbf{V} := (\text{Shift}(\mathbf{w}) \mid \text{Shift}(\mathbf{y}))$ of the sublattice $\mathcal{L}_{\mathbf{v}} := \mathcal{L}(\mathbf{V}) \subset \mathcal{L}_2$. Now, for r signatures $\mathbf{v}_1, \dots, \mathbf{v}_r$, these generating matrices $\mathbf{V}_1, \dots, \mathbf{V}_r$ generate the sublattice

$$\mathcal{L}_{\mathbf{v}_1, \dots, \mathbf{v}_r} := \mathcal{L}_{\mathbf{v}_1} + \dots + \mathcal{L}_{\mathbf{v}_r} := \left\{ \sum_{i=1}^r \mathbf{u}_i : \mathbf{u}_i \in \mathcal{L}_{\mathbf{v}_i} \quad \forall i \right\} \subset \mathcal{L}_2,$$

and for large enough r we can expect that in fact the equality $\mathcal{L}_{\mathbf{v}_1, \dots, \mathbf{v}_r} = \mathcal{L}_2$ holds. The Hermite normal form (HNF) is the analog of the echelon form for matrices with integer coefficients and can be used to efficiently extract a basis \mathbf{B}_r of $\mathcal{L}_{\mathbf{v}_1, \dots, \mathbf{v}_r}$ from all its generating vectors. More precisely, one has $\mathbf{B}_r = \text{HNF}((\mathbf{V}_1; \dots; \mathbf{V}_r))$, where the zero rows of the HNF are removed and the matrices \mathbf{V}_i are stacked vertically. To speed this process up in practice, it can be helpful to compute the HNF incrementally by adding $\mathbf{V}_1, \dots, \mathbf{V}_r$ step by step.

⁶ See the file `estimates/estimate_reduction.sage` available at <https://github.com/WvanWoerden/FuLeakage/>.

The remaining question is how large r has to be to guarantee the equality $\mathcal{L}_{v_1, \dots, v_r} = \mathcal{L}_2$ with overwhelming probability. Note that \mathcal{L}_2 is a rank- k lattice in a $2k$ -dimensional space, and that the rank remains the same if we restrict ourselves to the first k coefficients of each vector. In other words, we do not consider \mathcal{L}_2 generated by the quasi-circular shifts of \mathbf{g} anymore, but the lattice generated by $\text{Shift}(\mathbf{a})$. Recall from the preliminaries that this is in fact an ideal lattice $\mathcal{I} = (a) \subset \mathbb{Z}[x]/(x^k - 1)$ with generator $a = \sum_{i=1}^k a_i x^{i-1}$ corresponding to the secret vector $\mathbf{a} = (a_1, \dots, a_k)$. The question then becomes how many random elements we need to sample from \mathcal{I} such that they generate the full ideal.




Fig. 3: Number of signatures needed to fully generate the lattice \mathcal{L}_2 , leaving out signatures whose coefficients sum up to 0 (about 2.5% of the total).

This question has already been explored for an ideal over a number ring $\mathcal{O}_K \subset K$ and decent random distributions for sampling. In this scenario, the probability that r random elements fully generate the ideal is heuristically roughly $\zeta_K(r)^{-1}$, where ζ_K is the Dedekind zeta function of the number field K . See [16, Lemma 2.7] for a provable lower bound of $(4\zeta_K(r))^{-1}$. We have seen that $\mathbb{Z}[x]/(x^k - 1)$ splits into two number rings, a cyclotomic number ring $R_1 := \mathbb{Z}[x]/(\Psi_k)$ of degree $k - 1$, and a number ring $R_2 := \mathbb{Z}[x]/(x - 1) \cong \mathbb{Z}$ of degree 1. In order to generate an ideal $\mathcal{I} \subset \mathbb{Z}[x]/(x^k - 1)$, we need to generate its images $\mathcal{I}_1 := \mathcal{I} \pmod{\Psi_k}$ and $\mathcal{I}_2 := \mathcal{I} \pmod{x - 1}$ in R_1 and R_2 , respectively. Thus, we expect a probability of roughly $\zeta_{R_1}(r)^{-1} \cdot \zeta_{R_2}(r)^{-1}$. For prime cyclotomic number fields of large degree k , $\zeta_{R_1}(2)$ is close to 1 already and converges quickly for growing r . Concretely, we observe $\zeta_{R_1}(2) \leq 1.0024$ for primes $100 < k < 300$ and we obtain the upper bound $\zeta_{R_1}(r) \leq \exp(O(p^{\frac{3}{2}-r}))$ following the proof of [1, Lemma 1]. Thus, the probability is mostly dominated by $\zeta_{R_2}(r) = \zeta_{\mathbb{Z}}(r)$, which is the Riemann zeta function $\zeta(r)$. Given that $\zeta_{R_2}(r) \leq 1 + 3 \cdot 2^{-r}$ for $r \geq 2$, we can simply pick, say, $r \geq 100$ elements to guarantee that we generate the full ideal with all but negligible probability. We exclude the signatures that are $0 \in R_2$, i.e., those that have coefficients summing up to 0, as they can never help generating the ideal in R_2 , and such elements should be ignored for the Dedekind zeta estimate to hold.

In Figure 3, we verify the derived estimate experimentally over 1,000 trials. All ideals were successfully generated by at most 11 signatures.

To conclude, a few dozen signatures are enough to recover a basis for the lattice \mathcal{L}_2 containing the unusually short secret vectors $\mathbf{g}_1, \dots, \mathbf{g}_k$. Then, by the BKZ lattice-reduction algorithm, we can recover any such \mathbf{g}_i in time $2^{\frac{0.292k}{2} + o(k)}$, pushing the attack cost significantly below the claimed security levels as shown in Table 4.

This attack could be mitigated by adopting the basis-extension strategy from e.g. NTRUSIGN [20] and FALCON [31]. Namely, their signature generation does not only use the secret good basis \mathbf{G}_{sec} but its extension \mathbf{B}_{sec} that generates \mathcal{L}_1 . This makes sure that the generated signatures are distributed in the whole lattice \mathcal{L}_1 and do not only lie in the sublattice \mathcal{L}_2 .

5 A polynomial-time quantum attack

In this section, we present a polynomial-time quantum attack on all security levels of FuLeeca, hence breaking the post-quantum claims of the scheme. We show that the leaked lattice \mathcal{L}_2 from the previous section can be transformed into a related principal-ideal lattice over the cyclotomics with an unusually short generator linked to the secret key. For this family of ideals, there exists a polynomial-time quantum algorithm to recover the unusually short generator [7, 12]. Thus, we can recover the secret key in quantum-polynomial time.

5.1 Constructing a circulant ideal lattice

Recall that the lattice \mathcal{L}_2 defined in equation (2) is generated by the rows of $\mathbf{G}_{\text{sec}} = (\mathbf{A} \mid \mathbf{B})$ and that the blocks $\mathbf{A} = \text{Shift}(\mathbf{a})$ and $\mathbf{B} = \text{Shift}(\mathbf{b})$ are circulant matrices. Let us consider the lattice

$$\mathcal{L}_3 := \left\{ \sum_{i=1}^k z_i \mathbf{a}_i : z_i \in \mathbb{Z} \quad \forall i \right\} \subset \mathbb{R}^k$$

that is spanned by the rows $\mathbf{a}_i = \text{shift}^{i-1}(\mathbf{a})$ of \mathbf{A} . We can recover a basis of \mathcal{L}_3 from any basis of \mathcal{L}_2 simply by considering only the first half of each basis vector. Since this lattice is circulant, it can be represented by an ideal \mathcal{I} in the quotient ring $\mathbb{Z}[x]/(x^k - 1)$ with $k = n/2$. More precisely, we can identify any vector $(f_1, \dots, f_k) \in \mathcal{L}_3$ with the element $f(x) = \sum_{i=1}^k f_i x^{i-1} \in \mathbb{Z}[x]/(x^k - 1)$ by means of the coefficient embedding discussed in subsection 2.3. Recall that the circular shift of a lattice vector coincides precisely with the multiplication of its polynomial representation by x modulo $x^k - 1$. Further, \mathcal{L}_3 corresponds precisely to the principal ideal $(a) \subset \mathbb{Z}[x]/(x^k - 1)$ generated by $a(x) = \sum_{i=1}^k a_i x^{i-1}$, where $\mathbf{a} = (a_1, \dots, a_k)$ is the secret vector. As \mathcal{I} is closed under multiplication with x , note that in fact every $x^i a$ for $i = 0, \dots, k - 1$ generates \mathcal{I} and shares its short Euclidean norm with a . Thus, it is enough to recover one of these short generators to find a shift of the secret vector \mathbf{a} .

We will use a quantum algorithm to recover a short generator of a principal ideal in a cyclotomic number ring, but so far we are working over the quotient $\mathbb{Z}[x]/(x^k - 1)$. Therefore, we switch to the image of \mathcal{I} in a suitable, slightly smaller, cyclotomic number ring. Namely, choose the cyclotomic polynomial

$$\psi_k(x) := (x^k - 1)/(x - 1) = 1 + x + \dots + x^{k-1}$$

and consider the cyclotomic number ring $\mathcal{O}_k := \mathbb{Z}[x]/(\psi_k) \subset \mathbb{Q}[x]/(\psi_k)$, as discussed in subsection 2.3. Any element f of $\mathbb{Z}[x]/(x^k - 1)$ can be transferred to \mathcal{O}_k by the ring homomorphism $f \pmod{x^k - 1} \mapsto f \pmod{\psi_k}$. This mapping preserves the ideal structure of $\mathcal{I} \subset \mathbb{Z}[x]/(x^k - 1)$ and its image in \mathcal{O}_k is the principal ideal $\mathcal{I}' := \mathcal{I} \pmod{\psi_k}$ that is generated by $a \pmod{\psi_k}$, or, more generally, by any $x^i a \pmod{\psi_k}$ for $i = 0, \dots, k - 1$. In the case of cyclotomic rings, this transformation only changes the geometry slightly, i.e., polynomials with small coefficients in \mathcal{I} are generally still small in the ideal lattice \mathcal{I}' under the canonical embedding [8].

5.2 Quantumly recovering the secret key from the ideal lattice

We now exploit the leaked lattice \mathcal{L}_3 to fully recover the secret key in quantum-polynomial time. In particular, we make use of a quantum algorithm which can recover a shortest generator of any principal ideal within a cyclotomic number ring that has an exceptionally short generator [7, 12].

As we have seen in the previous subsection, we can identify \mathcal{L}_3 with a principal ideal \mathcal{I} in $\mathbb{Z}[x]/(x^k - 1)$ and any short generator $x^i a$ of \mathcal{I} corresponds to a shift of the secret vector \mathbf{a} . Furthermore, we can easily map \mathcal{I} to a principal ideal \mathcal{I}' in the cyclotomic number ring $\mathcal{O}_k = \mathbb{Z}[x]/(\psi_k)$. The strategy to recover a suitable short generator $x^i a$ of \mathcal{I} , and thus a shift of the secret vector \mathbf{a} , is to firstly find a principal generator c of \mathcal{I}' , secondly transform it into a *short* generator a' of \mathcal{I}' , and thirdly lift it back to a generator of \mathcal{I} in the quotient ring $\mathbb{Z}[x]/(x^k - 1)$. Figure 2 visualizes the relations between the involved lattices.

Recovering a principal generator of \mathcal{I}' . Since we can obtain a basis of \mathcal{L}_3 from any basis of \mathcal{L}_2 and the latter can be computed from a small sample of FuLeeca signatures as described in subsection 4.2, we have access to a set of generators c_1, \dots, c_s of \mathcal{I} . Their images $c_1 \pmod{\psi_k}, \dots, c_s \pmod{\psi_k}$ then span $\mathcal{I}' \subset \mathcal{O}_k$. Recovering a principal generator of \mathcal{I}' from these generators classically takes sub-exponential time, but there is a quantum algorithm achieving that in only polynomial quantum time [7]. We can thus efficiently obtain a generator $c \in \mathcal{I}'$ with $\mathcal{I}' = (c)$.

Shortening the principal generator of \mathcal{I}' . Next, we transform c into a short generator a' of \mathcal{I}' . Since any two generators of \mathcal{I}' differ by multiplication with a unit in \mathcal{O}_k , there is a unit $u \in \mathcal{O}_k^\times$ satisfying $c = u \cdot a'$. This equality reads as

$$\text{Log}(c) = \text{Log}(u) + \text{Log}(a')$$

in terms of the logarithmic embedding, which was discussed in subsection 2.3. Note that $\text{Log}(u)$ lies in the log-unit lattice $\text{Log}(\mathcal{O}_k^\times)$ by definition, while $\text{Log}(c)$ and $\text{Log}(a')$ are contained in the larger set $\text{Log}(K^\times)$ but in general not in $\text{Log}(\mathcal{O}_k^\times)$. In fact, $\text{Log}(c)$ and $\text{Log}(a')$ do not even need to lie in the real span of the log-unit lattice, which is contained in the subspace orthogonal to the all-one vector $\mathbf{1} = (1, \dots, 1) \in \mathbb{R}^{k-1}$. This originates from $\text{Log}(\mathcal{O}_k^\times)$ not being full-rank in $\text{Log}(K^\times)$.⁷

As we want to make use of decoding properties of the log-unit lattice, we consider the orthogonal projection $\pi : \mathbb{R}^{k-1} \rightarrow \text{span}(\text{Log}(\mathcal{O}_k^\times))$ onto the log-unit lattice and transform the above equality into

$$\pi(\text{Log}(c)) = \text{Log}(u) + \pi(\text{Log}(a')). \quad (3)$$

As illustrated in Figure 4, this can now be interpreted as a decoding problem in the log-unit lattice, where $\pi(\text{Log}(c))$ is the target, $\text{Log}(u)$ the desired lattice element, and $\pi(\text{Log}(a'))$ the error. If we solve this decoding problem and recover $\text{Log}(u)$, we can also compute the unit $u \in \mathcal{O}_k$ and thus get a short generator $a' = c/u$ of \mathcal{I}' , as desired.




Fig. 4: Visualization of the decoding problem $\pi(\text{Log}(c)) = \text{Log}(u) + \pi(\text{Log}(a'))$ after projecting onto the log-unit lattice.

Decoding in the log-unit lattice. Let us now explain why equation (3) indeed gives us a decoding problem in the log-unit lattice $\text{Log}(\mathcal{O}_k^\times)$ that we can

⁷ $\text{Log}(K^\times)$ is not full-rank in \mathbb{R}^{k-1} either, as the complex embeddings come in conjugate pairs. However, we are only concerned with the fact that $\text{Log}(\mathcal{O}_k^\times)$ is not full-rank in $\text{Log}(K^\times)$.

decode efficiently to the unique solution $\text{Log}(u)$. We will also provide experimental evidence that this is indeed the case for the **FuLecca** setting at the end of the section. For uniqueness, and later for efficient decoding, we need the error $\|\pi(\text{Log}(a'))\|_2$ to be *small*. Then the target $\pi(\text{Log}(c))$ lies exceptionally close to the log-unit lattice vector $\text{Log}(u)$, and far away from the rest of the log-unit lattice.

We argue intuitively why the error $\|\pi(\text{Log}(a'))\|_2$ is indeed small. Note first that the algebraic norm $N(a') := \prod_{i=1}^{k-1} |\sigma_i(a')|$ coincides for any principal generator of \mathcal{I}' , as every unit u satisfies $N(u) = 1$ and the norm is multiplicative. But at the same time, we assume that

$$\|a'\|_2^2 = \sum_{i=1}^{k-1} |\sigma_i(a')|^2$$

is exceptionally small in the canonical embedding because a' is a *short* generator of \mathcal{I}' . Thus, the summands $|\sigma_i(a')|^2$ for $i = 1, \dots, k-1$ need to be roughly balanced in terms of absolute size. But this holds true also for the logarithm of their square roots, and thus the sizes of the coefficients of $\text{Log}(a')$ are approximately equal.

Recall that the log-unit lattice $\text{Log}(\mathcal{O}_k^\times)$ lies in the hyperplane orthogonal to the vector $\mathbf{1} = (1, \dots, 1) \in \mathbb{R}^{k-1}$. Because we expect the logarithmic embedding $\text{Log}(a')$ to have approximately equal coefficients, the projection $\pi(\text{Log}(a'))$ is expected to be short. I.e., we expect $\pi(\text{Log}(c))$ to lie exceptionally close to the log-unit lattice, and the recovery of $\text{Log}(u)$ from $\pi(\text{Log}(c))$ is a unique decoding problem in the log-unit lattice.

For efficient decoding in the log-unit lattice, we make use of a good basis \mathbf{B} of $\text{Log}(\mathcal{O}_k)$ and its pseudo-inverse \mathbf{B}^\dagger . In particular, we choose \mathbf{B} as in [12], i.e., consisting of the units $(x^j - 1)/(x - 1)$ for $j = 2, \dots, (k-1)/2$.⁸ We can write $\pi(\text{Log}(c)) = \mathbf{y}\mathbf{B}$ for some non-integer vector $\mathbf{y} \in \mathbb{R}^{(k-3)/2}$, and thus $\mathbf{y} = \pi(\text{Log}(c)) \cdot \mathbf{B}^\dagger$ in terms of the pseudo-inverse \mathbf{B}^\dagger . Now because \mathbf{B} is a good basis and $\pi(\text{Log}(c))$ lies exceptionally close to the lattice, we can round the coefficients of \mathbf{y} to obtain the nearby lattice point

$$\text{Log}(u) = \lfloor \mathbf{y} \rfloor \mathbf{B} = \lfloor \pi(\text{Log}(c)) \cdot \mathbf{B}^\dagger \rfloor \mathbf{B}.$$

The precise condition for the above to be true is that

$$\pi(\text{Log}(c)) - \text{Log}(u) = \pi(\text{Log}(a')) \in \left(-\frac{1}{2}, \frac{1}{2}\right)^{(k-3)/2} \mathbf{B},$$

or, equivalently, that all coefficients of $\pi(\text{Log}(a')) \cdot \mathbf{B}^\dagger$ lie in the interval $(-\frac{1}{2}, \frac{1}{2})$, which is the case if $\pi(\text{Log}(a'))$ is small enough.

⁸ As discussed in [12], \mathbf{B} does *not* necessarily generate the full log-unit lattice but only a sublattice thereof. To still decode efficiently in the full lattice, we only need the index of the sublattice to be small, and that seems to be the case under reasonable heuristics. In particular, the index is less than 11 for all primes $k \leq 241$. If desired, a full basis for the log-unit lattice can efficiently be computed by a quantum algorithm.

The decoding process could be further improved by considering Babai’s nearest-plane decoding instead of simple coefficient rounding. However, this is not needed in our case.

Lifting the result and recovering the key. From $\text{Log}(u)$ we obtain the unit u up to some root of unity x^{-i} and therefore a short generator $c/(x^{-i}u) = x^i a'$ of \mathcal{I}' . Recall that the quotient ring $R = \mathbb{Z}[x]/(x^k - 1)$ splits into a product of number rings $R_1 = \mathbb{Z}[x]/(\psi_k)$ and $R_2 = \mathbb{Z}[x]/(x - 1)$. So to lift $a' = x^i a \pmod{\psi_k} \in R_1$ to $x^i a \in R$ we furthermore need to know the value of $x^i a \pmod{x - 1} = a \pmod{x - 1} \in R_2$. Note that $a \pmod{x - 1} = a(1)$. One can either guess this small value, or use that for any set c_1, \dots, c_s of generators for \mathcal{I} we have $a(1) = \text{gcd}(c_1(1), \dots, c_s(1))$.

To conclude, we obtain $x^i a$, and thus a shift $\text{shift}^i(\mathbf{a})$ of the secret vector \mathbf{a} . Using the public key, or alternatively a basis of \mathcal{L}_2 , we then receive a shift of \mathbf{b} , and thus have achieved a full key recovery.

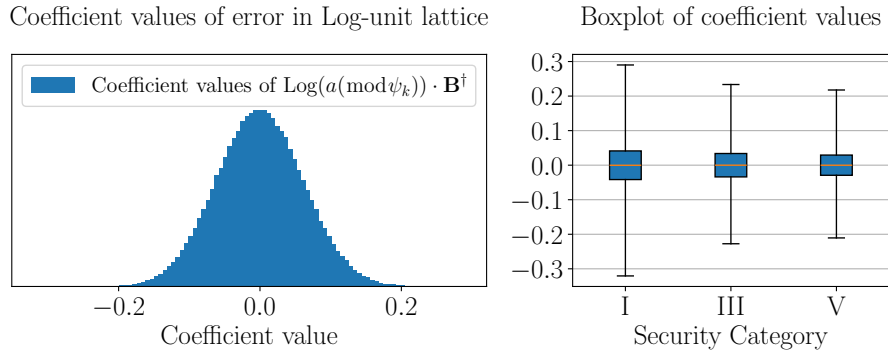


Fig. 5: On the left, distribution of the coefficients of $\text{Log}(a \pmod{\psi_k}) \cdot \mathbf{B}^\dagger$ for FuLeeca-I obtained from 1,000 samples. On the right, a boxplot of the same values for all parameter sets showing the standard deviation as well as minimum and maximum values.

Experimental confirmation. We have yet to verify that $\pi(\text{Log}(a'))$ is indeed small enough such that $\text{Log}(u)$ is efficiently recovered by the rounding procedure. To show that this is the case, we generated 1,000 secret-key instances for each security level of FuLeeca. For each instance, we confirmed that the coefficients of $\pi(\text{Log}(a')) \cdot \mathbf{B}^\dagger$ are indeed all well within the range of $(-\frac{1}{2}, \frac{1}{2})$. In fact, the $\frac{k-3}{2} \cdot 1,000$ coefficients of the 1,000 signature samples for the parameter sets FuLeeca-I, FuLeeca-III, and FuLeeca-V are centered around zero with a standard deviation of 0.0617, 0.0503, and 0.0435, and with a maximal absolute value of at

most 0.3206, 0.2335, and 0.2176, respectively. Figure 5 shows a full histogram of the distribution for FuLeeca-I and boxplots for all security levels.

Remark 1. For certain nice distributions of the unexpectedly short generator a' , one can *prove* analytically that the decoding in the log-unit lattice will succeed with high probability [12]. For FuLeeca however, the distribution seems too ad-hoc to easily prove a similar statement. From the perspective of an attacker, the numerical validation of this step is sufficient.

6 A full learning attack

In this section, we present a learning attack on FuLeeca that recovers the full secret key for all parameter sets FuLeeca-I, FuLeeca-III, and FuLeeca-V with only 90,000, 175,000, and 175,000 signatures, respectively. The attack exploits a bias introduced by the concentration procedure within the signing algorithm. Furthermore, we show a simple way to remove the bias and thus prevent this attack.

6.1 Observing and precomputing the bias

Recall that the signing procedure consists of a simple-signing step and a concentration step, which are depicted in Algorithm 2. The output is a signature \mathbf{v} that is the product $\mathbf{v} = \mathbf{x}\mathbf{G}_{\text{sec}}$ of a suitable vector $\mathbf{x} \in \mathbb{F}_p^k$ and the secret key $\mathbf{G}_{\text{sec}} \in \mathbb{F}_p^{k \times n}$. Essentially, the concentration step tries to concentrate the Lee weight of \mathbf{v} and the logarithmic matching probability $\text{LMP}(\mathbf{c}, \mathbf{v})$ around prescribed values, where $\mathbf{c} \in \{\pm 1\}^n$ is the message hash. Recall that the value of $\text{LMP}(\mathbf{c}, \mathbf{v})$ is directly determined by $\text{mt}(\mathbf{c}, \mathbf{v})$. This number of sign matches $\text{mt}(\mathbf{c}, \mathbf{v})$ of \mathbf{c} and \mathbf{v} is concentrated by successively adding or subtracting the rows $\mathbf{g}_1, \dots, \mathbf{g}_k$ of the secret key \mathbf{G}_{sec} . The procedure is similar to the iterative slicing algorithm for finding close lattice vectors [36].

The algorithm proceeds by trying to improve the number of sign matches first by adding or subtracting \mathbf{g}_1 , then \mathbf{g}_2 , all the way up to \mathbf{g}_k . However, the order in which this is done introduces a significant bias. More precisely, the first vector \mathbf{g}_1 is considered first every time and thus has a higher probability to be added or subtracted in contrast to, say, the last vector \mathbf{g}_k . This is visualized in Figure 6b for two dimensions, where the distribution of the signatures rather creates an elliptic shape than a circle and thus clearly deviates from the expected outcome in a non-biased case.

Recall that adding $\pm \mathbf{g}_i$ to a signature $\mathbf{v} = \mathbf{x}\mathbf{G}_{\text{sec}}$ corresponds to increasing or decreasing the i -th coordinate x_i of \mathbf{x} for $i = 1, \dots, k$. Indeed, we observe for the signatures generated by the reference implementation of FuLeeca that the first coefficients of \mathbf{x} are smaller than the latter ones. More concretely, we observe that the average values $\text{Avg}[x_1^2], \dots, \text{Avg}[x_k^2]$ are increasing, as shown in Figure 6a. We use an approximation of these biased average values for the learning attack. Therefore, we precomputed a single key and 2,500,000 signatures

for each parameter setting, and saved the average values as $\tilde{\mathbf{d}} = (\tilde{d}_1, \dots, \tilde{d}_k) := (\text{Avg}[x_1^2], \dots, \text{Avg}[x_k^2])$. Figure 6a shows that not all keys have exactly the same bias, but their individual values turn out to be sufficiently close to the precomputed data to permit our attack.




Fig. 6: Observed bias for FuLeeca signatures $\mathbf{v} = \mathbf{x}\mathbf{G}_{\text{sec}}$. On the left, experimental observations of the average values $\text{Avg}[x_i^2]$ for $i = 1, \dots, k$ for all parameter sets. The curves depict the results for 2,500,000 signatures from one key, i.e., the precomputed values $\tilde{d}_1, \dots, \tilde{d}_k$ used in the learning attack. The fog around the curves corresponds to 100,000 signatures generated for each of 50 random keys. The data points around the dashed lines show the result from one key and 100,000 signatures after properly randomizing the ordering in the concentration procedure. On the right, a two-dimensional visualization of the signature bias that clearly deviates from a circle.

To confirm that the fixed ordering of the concentration procedure is the cause of this bias, we adapted the reference implementation of FuLeeca to properly randomize the order in each iteration of the concentration procedure, and depicted the outcome in Figure 6a. Indeed, the previously biased average values $\text{Avg}[x_1^2], \dots, \text{Avg}[x_k^2]$ are roughly constant after this adaptation. This simple solution prevents the learning attack that follows. However, as long as there is no proof of non-leakage we do not expect this solution to hold against other learning attacks. In other schemes like FALCON [31] and Wave [5] this is achieved by following the GPV framework [17]. How to adapt these methods and proofs to the Lee metric is a non-trivial question, but one we think is necessary to answer in order to prevent other learning attacks on FuLeeca.

Remark 2. The learning attack on NTRUSIGN [27] uses the quasi-cyclic structure to turn a single signature sample into k distinct signatures by considering

all quasi-circular shifts. Thereby, the number of needed signature vectors for the attack is severely reduced to a few hundred. This is not possible in our setting because the observed bias depends precisely on the fact that the signature distribution is *not* invariant under quasi-circular shifts.

6.2 Approximating the secret key from the bias

We now exploit the observed bias to recover an approximation of the secret vector $\mathbf{g} = (\mathbf{a} \mid \mathbf{b})$. In fact, we will first retrieve its first half $\pm\mathbf{a}$ up to sign and from this recover $\pm\mathbf{b}$ as well. Note that $-\mathbf{a}$ and $-\mathbf{b}$ work just as well for generating signatures and can be considered as an alternative secret key. The attack only uses the first half $\mathbf{w} \in \mathbb{Z}^k$ of each signature vector $\mathbf{v} = \mathbf{x}\mathbf{G}_{\text{sec}} \in \mathbb{Z}^n$ and the bias that we obtain therein. We call \mathbf{w} a *partial signature* for simplicity and remark that it can be expressed as $\mathbf{w} = \mathbf{x}\mathbf{A}$ because of the block structure of $\mathbf{G}_{\text{sec}} = (\mathbf{A} \mid \mathbf{B})$ with $\mathbf{A} = \text{Shift}(\mathbf{a})$ and $\mathbf{B} = \text{Shift}(\mathbf{b})$.

Remark 3. The attack can easily be generalized to use the full signatures, and as such directly recover \mathbf{a} and \mathbf{b} . However, this does not seem to significantly improve the performance. In certain cases, it even decreases the success rate, while leading to higher computational costs because of the larger vectors.

Learning from expectations. Assume that N partial signatures $\mathbf{w}_1, \dots, \mathbf{w}_N$ are available. We can compute the average outer product

$$\text{Avg}[\mathbf{w}^\top \mathbf{w}] := \frac{1}{N} \sum_{i=1}^N \mathbf{w}_i^\top \mathbf{w}_i \approx \mathbb{E}[\mathbf{w}^\top \mathbf{w}]$$

to approximate the expectation of the outer product $\mathbf{w}^\top \mathbf{w} \in \mathbb{R}^{k \times k}$. Clearly, the quality of this estimate depends on the number N of given signature vectors. We will use this quantity to recover an approximation of the secret vector $\pm\mathbf{a}$.

Formally, we can rewrite the expected value of $\mathbf{w}^\top \mathbf{w}$ for $\mathbf{w} = \mathbf{x}\mathbf{A}$ in terms of the expected value of $\mathbf{x}^\top \mathbf{x}$ as

$$\mathbb{E}[\mathbf{w}^\top \mathbf{w}] = \mathbb{E}[\mathbf{A}^\top \mathbf{x}^\top \mathbf{x} \mathbf{A}] = \mathbf{A}^\top \cdot \mathbb{E}[\mathbf{x}^\top \mathbf{x}] \cdot \mathbf{A} = \mathbf{A}^\top \mathbf{D} \mathbf{A} + \mathbf{A}^\top \mathbf{R} \mathbf{A},$$

where the last step decomposes the expectation $\mathbb{E}[\mathbf{x}^\top \mathbf{x}]$ into its diagonal part $\mathbf{D} = \text{diag}(d_1, \dots, d_k)$ with $d_i = \mathbb{E}[x_i^2]$, and the remainder $\mathbf{R} = \mathbb{E}[\mathbf{x}^\top \mathbf{x}] - \mathbf{D}$ with zero diagonal. We write

$$\text{Avg}[\mathbf{w}^\top \mathbf{w}] = \mathbf{A}^\top \mathbf{D} \mathbf{A} + \mathbf{E} \quad \text{with} \quad \mathbf{D} = \text{diag}(d_1, \dots, d_k), \quad (4)$$

where \mathbf{E} is an error matrix that includes $\mathbf{A}^\top \mathbf{R} \mathbf{A}$ and the error introduced by the approximation $\text{Avg}[\mathbf{w}^\top \mathbf{w}] \approx \mathbb{E}[\mathbf{w}^\top \mathbf{w}]$. We discuss in the following how $\pm\mathbf{a}$ can be recovered if no error \mathbf{E} is present and then generalize the approach to the erroneous scenario.

Recovering $\pm \mathbf{a}$ when $\mathbf{E} = \mathbf{0}$. Let us for simplicity start with the error-free case, that is, with $\mathbf{E} = \mathbf{0}$ and thus $\text{Avg}[\mathbf{w}^\top \mathbf{w}] = \mathbf{A}^\top \mathbf{D} \mathbf{A}$. The following lemma shows that we can fully recover $\pm \mathbf{a}$ as soon as we know \mathbf{D} or, equivalently, its diagonal entries d_1, \dots, d_k .

Lemma 1. *Let $\mathbf{d} = (d_1, \dots, d_k) \in \mathbb{R}^k$ be a vector for which the matrix $\text{Shift}(\mathbf{d})$ is invertible. Write $\mathbf{D} = \text{diag}(\mathbf{d})$ and let $\mathbf{A} = \text{Shift}(\mathbf{a})$ be circulant for a vector $\mathbf{a} \in \mathbb{R}^k$. Then, $\pm \mathbf{a}$ can be efficiently recovered from \mathbf{d} and $\mathbf{A}^\top \mathbf{D} \mathbf{A}$.*

Proof. Note that the i -th row of the matrix $\mathbf{D} \mathbf{A}$ is given by $d_i \cdot \text{shift}^{i-1}(\mathbf{a})$ for $i = 1, \dots, k$. Therefore, we can express $\mathbf{A}^\top \mathbf{D} \mathbf{A}$ as

$$\mathbf{A}^\top \mathbf{D} \mathbf{A} = \sum_{i=1}^k d_i \cdot (\text{shift}^{i-1}(\mathbf{a}))^\top \text{shift}^{i-1}(\mathbf{a}),$$

i.e., as a weighted sum of outer products. Now fix some $j = 1, \dots, k$ and consider the j -th diagonal of the above matrix equation. Since shifting a vector and taking the j -th diagonal of its outer product commute, we obtain

$$\begin{aligned} \text{diag}_j(\mathbf{A}^\top \mathbf{D} \mathbf{A}) &= \sum_{i=1}^k d_i \cdot \text{diag}_j\left(\left(\text{shift}^{i-1}(\mathbf{a})\right)^\top \text{shift}^{i-1}(\mathbf{a})\right) \\ &= \sum_{i=1}^k d_i \cdot \text{shift}^{i-1}\left(\text{diag}_j(\mathbf{a}^\top \mathbf{a})\right) \\ &= \mathbf{d} \cdot \text{Shift}\left(\text{diag}_j(\mathbf{a}^\top \mathbf{a})\right). \end{aligned}$$

Recall that $\mathbf{x} \cdot \text{Shift}(\mathbf{y}) = \mathbf{y} \cdot \text{Shift}(\mathbf{x})$ holds for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^k$ to get

$$\text{diag}_j(\mathbf{A}^\top \mathbf{D} \mathbf{A}) = \mathbf{d} \cdot \text{Shift}\left(\text{diag}_j(\mathbf{a}^\top \mathbf{a})\right) = \text{diag}_j(\mathbf{a}^\top \mathbf{a}) \cdot \text{Shift}(\mathbf{d}).$$

Because $\text{Shift}(\mathbf{d})$ is invertible, we can solve the obtained linear system to recover $\text{diag}_j(\mathbf{a}^\top \mathbf{a})$. After repeating the above for all diagonals $j = 1, \dots, k$, we thus recover the outer product $\mathbf{a}^\top \mathbf{a}$, from which $\pm \mathbf{a}$ is easily computed. \square

Note that we can apply Lemma 1 to recover $\pm \mathbf{a}$ in the error-free FuLeeca setting, because the observed bias makes the vector \mathbf{d} strictly increasing and therefore the matrix $\text{Shift}(\mathbf{d})$ invertible. For a constant vector \mathbf{d} , the matrix $\text{Shift}(\mathbf{d})$ would only have rank 1 and the above recovery strategy would fail.

Approximating $\pm \mathbf{a}$ when $\mathbf{E} \neq \mathbf{0}$. Let us now consider the more general case in which \mathbf{d} is unknown and $\mathbf{E} \neq \mathbf{0}$ applies. We will use our precomputed approximation $\tilde{\mathbf{d}}$ of \mathbf{d} from subsection 6.1 to obtain an approximation $\tilde{\mathbf{a}}$ of $\pm \mathbf{a}$ using the same steps as in Lemma 1.

In the proof of Lemma 1, the entries of $\mathbf{a}^\top \mathbf{a}$ get computed from a linear system with coefficient matrix $\text{Shift}(\mathbf{d})$. The same computation with the approximation $\text{Shift}(\tilde{\mathbf{d}})$ therefore leads to an approximation

$$\mathbf{Y} = \mathbf{a}^\top \mathbf{a} + \mathbf{E}_d$$

of $\mathbf{a}^\top \mathbf{a}$, where \mathbf{E}_d is an error matrix that depends on the quality of the estimate $\tilde{\mathbf{d}}$ of \mathbf{d} . Because our approximation $\tilde{\mathbf{d}}$ of \mathbf{d} is good as shown in Figure 6a, we expect \mathbf{E}_d to be quite small.

In contrast to the setting of Lemma 1, we only know the erroneous version of $\mathbf{A}^\top \mathbf{D} \mathbf{A}$ displayed in equation (4). Thus, we now focus on the error matrix \mathbf{E} that is the sum of $\mathbf{A}^\top \mathbf{R} \mathbf{A}$ and the error introduced by the approximation $\text{Avg}[\mathbf{w}^\top \mathbf{w}] \approx \mathbb{E}[\mathbf{w}^\top \mathbf{w}]$. As we expect the approximation error to be small for a large enough number of signatures, \mathbf{E} is eventually dominated by $\mathbf{A}^\top \mathbf{R} \mathbf{A}$. Recall that for simple signing we have $\mathbf{x} = \text{trunc}(s/2 \cdot \mathbf{c} \cdot \text{sgn}(\mathbf{G}_{\text{sec}}^\top))$, and by the quasi-circulant structure of \mathbf{G}_{sec} the expectations $\mathbb{E}[\mathbf{x}^\top \mathbf{x}]$ and $\mathbb{E}[\mathbf{A}^\top \mathbf{x}^\top \mathbf{x} \mathbf{A}]$ are circulant, i.e. they have constant diagonals. After the concentration procedure, experiments show that all diagonals of \mathbf{E} are still roughly constant, i.e., that $\mathbf{E} \approx \text{Shift}(\mathbf{e})$ applies for some error vector $\mathbf{e} = (e_1, \dots, e_k) \in \mathbb{R}^k$. For $j = 1, \dots, k$, we thus get

$$\begin{aligned} \text{diag}_j(\text{Avg}[\mathbf{w}^\top \mathbf{w}]) &\approx \text{diag}_j(\mathbf{A}^\top \mathbf{D} \mathbf{A} + \text{Shift}(\mathbf{e})) \\ &= \text{diag}_j(\mathbf{a}^\top \mathbf{a}) \cdot \text{Shift}(\mathbf{d}) + (e_j, \dots, e_j) \end{aligned}$$

and, since $(1, \dots, 1)$ is an eigenvector of the circulant matrix $\text{Shift}(\mathbf{d})$,

$$\text{diag}_j(\text{Avg}[\mathbf{w}^\top \mathbf{w}]) \cdot \text{Shift}(\mathbf{d})^{-1} \approx \text{diag}_j(\mathbf{a}^\top \mathbf{a}) + \frac{1}{\lambda} \cdot (e_j, \dots, e_j) \quad (5)$$

with $\lambda \in \mathbb{R}$ being the corresponding eigenvalue. Because the entries $d_i = \mathbb{E}[x_i^2]$ of \mathbf{d} are positive and thus also $\text{Shift}(\mathbf{d})$ only contains positive entries, the eigenvalue $\lambda = \sum_{i=1}^k d_i$ is very large compared to the other eigenvalues. Experimentally, λ is between 9 and 18 times larger than the second biggest eigenvalue, and between 200 and 600 times larger than the average eigenvalue.

As a result, the error $e'_j := e_j/\lambda$ is reduced a lot and becomes relatively small, while in contrast $\text{diag}_j(\mathbf{a}^\top \mathbf{a})$ does not get scaled down as much because it is expected to have a somewhat random direction. Overall, we thus obtain an approximation $\mathbf{Y} \approx \mathbf{a}^\top \mathbf{a} + \text{Shift}(\mathbf{e}')$ of $\mathbf{a}^\top \mathbf{a}$ from equation (5), where $\mathbf{e}' := (e'_1, \dots, e'_k)$ is relatively small.

To conclude, we note that \mathbf{Y} is approximately a sum of a rank-1 matrix $\mathbf{a}^\top \mathbf{a}$ and a circulant error matrix $\text{Shift}(\mathbf{e}')$ with relatively small entries. We can thus hope to recover an approximation $\tilde{\mathbf{a}}$ of $\pm \mathbf{a}$ by computing an optimal rank-1 approximation of \mathbf{Y} . The latter is given by $\tilde{\mathbf{a}} = \sqrt{\sigma_1} \mathbf{u}_1$, where σ_1 and \mathbf{u}_1 are obtained from the singular-value decomposition

$$\mathbf{Y} = \sum_{i=1}^k \sigma_i \mathbf{u}_i^\top \mathbf{u}_i$$

with $\mathbf{u}_1, \dots, \mathbf{u}_k \in \mathbb{R}^k$ and the singular values $\sigma_1, \dots, \sigma_k \in \mathbb{R}$. Note that $\sigma_1 \geq \dots \geq \sigma_k$ holds and that the decomposition can be computed efficiently. For a low number of signatures, the error contribution can still be large and thus one might want to consider $\tilde{\mathbf{a}} = \sqrt{\sigma_i} \mathbf{u}_i$ for some small i . In the full attack, we simply

Algorithm 3: Approximate secret-key recovery.

Input : Partial FuLeeca signatures $\mathbf{w}_1, \dots, \mathbf{w}_N$,
 approximation $\tilde{\mathbf{d}}$ of $\mathbf{d} = (\mathbb{E}[x_1^2], \dots, \mathbb{E}[x_k^2])$.
Output : Approximation $\tilde{\mathbf{a}}$ of $\pm \mathbf{a}$.

```

1  $\mathbf{W} = \frac{1}{N} \sum_{i=1}^N \mathbf{w}_i^\top \mathbf{w}_i$ 
2  $\mathbf{Y} = \mathbf{0} \in \mathbb{R}^{k \times k}$ 
3 for  $j = 1, \dots, k$  do
4    $\text{diag}_j(\mathbf{Y}) = \text{diag}_j(\mathbf{W}) \cdot \text{Shift}(\tilde{\mathbf{d}})^{-1}$ 
5 end
6 return  $\text{argmin}_{\tilde{\mathbf{a}} \in \mathbb{R}^k} (\mathbf{Y} - \tilde{\mathbf{a}}^\top \tilde{\mathbf{a}})$  // best rank-1 approximation

```

start with $i = 1$ and continue with larger values until the secret key is recovered or we abort. Algorithm 3 summarizes the resulting procedure.

Figure 7 shows that the distance between $\tilde{\mathbf{a}}$ and $\pm \mathbf{a}$ is indeed small and decreases when the number of signatures increases, i.e., when the approximation $\text{Avg}[\mathbf{w}^\top \mathbf{w}] \approx \mathbb{E}[\mathbf{w}^\top \mathbf{w}]$ improves. For a large number of signatures, the distance stays roughly the same, reflecting the fixed errors \mathbf{E}_d and $\text{Shift}(\mathbf{e}')$ which are independent of the number of signatures.




Fig. 7: Relative distance of approximate solution $\tilde{\mathbf{a}}$ to $\pm \mathbf{a}$ after Algorithm 3 and after one averaging iteration, respectively. Further, success rates for recovering $\pm \mathbf{a}$ by the exact solving procedure from subsection 6.3 with the approximations obtained after no or one averaging step, respectively.

6.3 Recovering the secret key from its approximation

Given an approximation $\tilde{\mathbf{a}}$ of the secret vector $\pm\mathbf{a}$, we now try to recover the exact solution. As $\tilde{\mathbf{a}}$ already determines whether we recover \mathbf{a} or $-\mathbf{a}$ and both options are valid alternative secret keys, we stick to \mathbf{a} for brevity. We start with the general idea and then discuss further improvements that generally increase the success rate and therefore require less signatures to succeed.

General idea. The general idea is to use the short partial signature vectors $\mathbf{w} = \mathbf{x}\mathbf{A}$ and the approximation $\tilde{\mathbf{A}} = \text{Shift}(\tilde{\mathbf{a}})$ of $\mathbf{A} = \text{Shift}(\mathbf{a})$, that was obtained as described in subsection 6.2. Given \mathbf{w} and $\tilde{\mathbf{A}}$, we can compute an approximation $\tilde{\mathbf{x}} = \tilde{\mathbf{A}}^{-1} \cdot \mathbf{w}$ of \mathbf{x} . More concretely, if $\tilde{\mathbf{A}}^{-1} = \mathbf{A}^{-1} + \mathbf{E}'$ holds for some small error matrix \mathbf{E}' , we obtain

$$\tilde{\mathbf{x}} - \mathbf{x} = (\mathbf{A}^{-1}\mathbf{w} + \mathbf{E}'\mathbf{w}) - \mathbf{x} = (\mathbf{x} - \mathbf{x}) + \mathbf{E}'\mathbf{w} = \mathbf{E}'\mathbf{w}.$$

Because both \mathbf{E}' and \mathbf{w} are small, $\tilde{\mathbf{x}}$ is close to \mathbf{x} . Additionally, we know that \mathbf{x} is an integral vector. We can thus recover $\mathbf{x} = \lfloor \tilde{\mathbf{x}} \rfloor$ by rounding the individual coefficients if $\|\mathbf{x} - \tilde{\mathbf{x}}\|_\infty < \frac{1}{2}$ applies. Once the correct \mathbf{x} is recovered, the equation $\mathbf{w} = \mathbf{x}\mathbf{A}$ allows to derive the circulant matrix \mathbf{A} and thus the secret vector \mathbf{a} . To increase the success rate, we can try the above for every available partial signature vector \mathbf{w} until we correctly recover \mathbf{a} .


We can detect whether or not we have successfully recovered the secret key \mathbf{a} as follows: If our guess of \mathbf{x} is correct, then we obtain the correct integer vector \mathbf{a} from the equation $\mathbf{w} = \mathbf{x}\mathbf{A}$. If our guess of \mathbf{x} was not correct, we would not expect the solution to be close-to-integral and have typical norm $\|\mathbf{a}\|_2$, and we use this to efficiently reject the failed trials. In practice, this approach perfectly distinguishes the correct from the incorrect cases.

For the cases that do pass the mentioned checks, we have several options to provably verify their correctness. One option is to compute a basis for the lattice $\mathcal{L}_3 = \mathcal{L}(\mathbf{A})$ generated by \mathbf{A} from the available signatures as in section 5. We can then check if our recovered vector belongs to this lattice and is short enough in the Euclidean metric. Another option is to compute $\mathbf{b} = \mathbf{a}\mathbf{A}^{-1}\mathbf{B} \pmod{p}$ using the public key $\mathbf{G}_{\text{pub}} = (\mathbf{I}_k \mid \mathbf{A}^{-1}\mathbf{B})$ and to verify that \mathbf{a} and \mathbf{b} are short enough in the Lee metric.

Making the vector typical. Due to the use of a typical vector in the key generation, we already know \mathbf{a} up to a signed permutation. We can thus try to round our approximation $\tilde{\mathbf{a}}$ to the closest typical vector. We do this by sorting the coefficients of $\tilde{\mathbf{a}}$ and of the typical vector by absolute value, respectively, and matching the results up. For example, the largest absolute value x of $\tilde{\mathbf{a}}$ gets transformed into $\text{sgn}(x) \cdot y$, where y is the largest absolute value in the typical vector. The operation is efficient and we have observed that this generally improves the distance $\|\mathbf{a} - \tilde{\mathbf{a}}\|_2$ compared to simply rounding $\tilde{\mathbf{a}}$.

Iterative solving and averaging over guesses. Instead of directly recovering the exact secret key \mathbf{a} , we can try to improve the approximation iteratively until the approximation is good enough to recover the secret key exactly. One approach to achieve this is averaging over the approximated solutions we obtain for \mathbf{a} , in the hope of balancing out the errors. More precisely, we obtain an approximate solution $[\tilde{\mathbf{x}}]$ of \mathbf{x} for any signature $\mathbf{w} = \mathbf{x}\mathbf{A}$ from the general method presented above, and thus we get an approximation \mathbf{a}' of \mathbf{a} for every signature. Averaging over all these solutions \mathbf{a}' yields a new approximate solution $\tilde{\mathbf{a}}$ which is hopefully better than our initial approximation.

We can iteratively continue the process until the approximation is good enough to directly recover the secret key. The iterative improvements are showcased in Figure 8a. Because the same signature equations are used continuously, we observed that rounding or making $\tilde{\mathbf{a}}$ typical after each iteration is important to not get stuck in a fixed point.



(a) Approximation $\tilde{\mathbf{a}}$ vs secret key \mathbf{a} . (b) Distance $\|\mathbf{a} - \tilde{\mathbf{a}}\|_2$ vs averaging iterations.

Fig. 8: Demonstration of first approximation guess and further improvements by iterative solving and averaging. The data has been collected using 40,000 signatures for each of 50 random FuLeeca-I keys.

Selecting near-integer entries. In the initial method, we try to recover \mathbf{a} exactly from a single signature $\mathbf{w} = \mathbf{x}\mathbf{A}$, hoping that all coefficients of $\tilde{\mathbf{x}}$ round correctly to those of \mathbf{x} . Note however, that coefficients of $\tilde{\mathbf{x}}$ that are already close to being integral have a higher chance of rounding correctly than those close to $\frac{1}{2} + \mathbb{Z}$. Additionally, each correctly rounded coefficient gives a linear equation on \mathbf{a} . We can thus consider multiple signatures and only select k coefficients that are exceptionally close-to-integral already, and then recover \mathbf{a} from the resulting linear system. This method is similar to the threshold rounding in [23, 30]. For simplicity, we did not use this idea for the final attack script, as it only seemed to

improve the required number of averaging iterations but not the overall success rate.

6.4 Running the learning attack

To demonstrate the attack, we generated 50 random keys and 100,000, 200,000, and 200,000 signatures for FuLeeca-I, FuLeeca-III, and FuLeeca-V parameters, respectively. We then ran the attack using the first N signatures for different values of N . The overall success rate of the attack versus the number of signatures is demonstrated in Figure 1. All keys get recovered with only 90,000, 175,000, and 175,000 signatures for the different security levels, respectively. More than 80% of the keys get recovered with only 40,000, 125,000, and 100,000 signatures, respectively. Note in particular that FuLeeca-V requires less signatures to break than FuLeeca-III. A possible explanation for this is that the values $\text{Avg}[x_i^2]$ seem to have less variance between distinct keys for FuLeeca-V and are thus closer to the precomputed values, leading to a smaller error in the computation.

Figure 7 shows that the averaging strategy to improve the approximation $\tilde{\mathbf{a}}$ of \mathbf{a} is highly effective and increases the success rate of recovering the exact secret key significantly. In Figure 8b, we consider the regime with a relatively low number of signatures. The multiple averaging iterations combined with making the approximation typical can still recover the full secret even if the initial approximation was not so good.

The average of the best running time for each key, using the optimal number of signatures, is about 2, 12, and 14 minutes on a single core, respectively. The fastest successful attack took only 56 seconds, while the longest successful attack took a total of 98 minutes. All running times were a small fraction of the time needed to generate the signatures. To conclude, we can efficiently break all FuLeeca parameter sets given less than 175,000 signatures.

Running the attack scripts. The source code of our learning attack is available and open-source.⁹ The main attack script is `attack/full_attack.py`. We provide pre-generated FuLeeca signatures for one fixed key per FuLeeca parameter set and a convenient shell script to run the three corresponding attack instances. First, execute `attack/1_download_sigs.sh` to download the signature samples. Then, run `attack/2_run_attack.sh` to start the learning attack and see how instances of all FuLeeca parameter sets are broken in real time.

Acknowledgments. Experiments presented in this paper were carried out using the PlaFRIM experimental testbed, supported by Inria, CNRS (LABRI and IMB), Université de Bordeaux, Bordeaux INP and Conseil Régional d’Aquitaine. W. van Woerden was supported by the CHARM ANR-NSF grant (ANR-21-CE94-0003).

⁹ See <https://github.com/WvanWoerden/FuLeakage/> for the code and instructions.

References

1. Albrecht, M., Bai, S., Ducas, L.: A subfield lattice attack on overstretched NTRU assumptions: Cryptanalysis of some FHE and graded encoding schemes. Cryptology ePrint Archive, Report 2016/127 (2016), <https://eprint.iacr.org/2016/127>
2. Albrecht, M., Ducas, L.: Lattice attacks on ntru and lwe: a history of refinements. Cryptology ePrint Archive (2021)
3. Albrecht, M., Ducas, L.: Lattice attacks on NTRU and LWE: A history of refinements. Cryptology ePrint Archive, Report 2021/799 (2021), <https://eprint.iacr.org/2021/799>
4. Albrecht, M.R., Player, R., Scott, S.: On the concrete hardness of learning with errors. *Journal of Mathematical Cryptology* **9**(3), 169–203 (2015). <https://doi.org/10.1515/jmc-2015-0016>
5. Banegas, G., Carrier, K., Chailloux, A., Couvreur, A., Debris-Alazard, T., Gaborit, P., Karpman, P., Loyer, J., Niederhagen, R., Sendrier, N., Smith, B., Tillich, J.P.: Wave. Tech. rep., National Institute of Standards and Technology (2023), <https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/round-1/spec-files/wave-spec-web.pdf>
6. Bariffi, J., Bartz, H., Liva, G., Rosenthal, J.: On the properties of error patterns in the constant Lee weight channel. In: International Zurich Seminar on Information and Communication (IZS 2022). Proceedings. pp. 44–48. ETH Zurich (2022). <https://doi.org/10.3929/ETHZ-B-000535277>
7. Biasse, J.F., Song, F.: Efficient quantum algorithms for computing class groups and solving the principal ideal problem in arbitrary degree number fields. In: Krauthgamer, R. (ed.) 27th SODA. pp. 893–902. ACM-SIAM, Arlington, VA, USA (Jan 10–12, 2016). <https://doi.org/10.1137/1.9781611974331.ch64>
8. Blanco-Chacón, I.: On the RLWE/PLWE equivalence for cyclotomic number fields. *Applicable Algebra in Engineering, Communication and Computing* **33**(1), 53–71 (2022). <https://doi.org/10.1007/S00200-020-00433-Z>
9. Bos, J.W., Bronchain, O., Ducas, L., Fehr, S., Huang, Y., Pornin, T., Postlethwaite, E.W., Prest, T., Pulles, L.N., van Woerden, W.: HAWK. Tech. rep., National Institute of Standards and Technology (2023), <https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/round-1/spec-files/hawk-spec-web.pdf>
10. Conway, J.H., Sloane, N.J.A.: Sphere packings, lattices and groups, vol. 290. Springer Science & Business Media (2013)
11. Courtois, N., Finiasz, M., Sendrier, N.: How to achieve a McEliece-based digital signature scheme. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 157–174. Springer, Heidelberg, Germany, Gold Coast, Australia (Dec 9–13, 2001). https://doi.org/10.1007/3-540-45682-1_10
12. Cramer, R., Ducas, L., Peikert, C., Regev, O.: Recovering short generators of principal ideals in cyclotomic rings. In: Fischlin, M., Coron, J.S. (eds.) EUROCRYPT 2016, Part II. LNCS, vol. 9666, pp. 559–585. Springer, Heidelberg, Germany, Vienna, Austria (May 8–12, 2016). https://doi.org/10.1007/978-3-662-49896-5_20
13. Dachman-Soled, D., Ducas, L., Gong, H., Rossi, M.: LWE with side information: Attacks and concrete security estimation. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020, Part II. LNCS, vol. 12171, pp. 329–358. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 17–21, 2020). https://doi.org/10.1007/978-3-030-56880-1_12

14. Ducas, L., Nguyen, P.Q.: Learning a zonotope and more: Cryptanalysis of NTRUSign countermeasures. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 433–450. Springer, Heidelberg, Germany, Beijing, China (Dec 2–6, 2012). https://doi.org/10.1007/978-3-642-34961-4_27
15. Ducas, L., Postlethwaite, E.W., Pulles, L.N., van Woerden, W.P.J.: Hawk: Module LIP makes lattice signatures fast, compact and simple. In: Agrawal, S., Lin, D. (eds.) ASIACRYPT 2022, Part IV. LNCS, vol. 13794, pp. 65–94. Springer, Heidelberg, Germany, Taipei, Taiwan (Dec 5–9, 2022). https://doi.org/10.1007/978-3-031-22972-5_3
16. Felderhoff, J., Pellet-Mary, A., Stehlé, D.: On module unique-SVP and NTRU. In: Agrawal, S., Lin, D. (eds.) ASIACRYPT 2022, Part III. LNCS, vol. 13793, pp. 709–740. Springer, Heidelberg, Germany, Taipei, Taiwan (Dec 5–9, 2022). https://doi.org/10.1007/978-3-031-22969-5_24
17. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Ladner, R.E., Dwork, C. (eds.) 40th ACM STOC. pp. 197–206. ACM Press, Victoria, BC, Canada (May 17–20, 2008). <https://doi.org/10.1145/1374376.1374407>
18. 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, Heidelberg, Germany, Amsterdam, The Netherlands (Apr 28 – May 2, 2002). https://doi.org/10.1007/3-540-46035-7_20
19. Goldreich, O., Goldwasser, S., Halevi, S.: Public-key cryptosystems from lattice reduction problems. In: Kaliski Jr., B.S. (ed.) CRYPTO’97. LNCS, vol. 1294, pp. 112–131. Springer, Heidelberg, Germany, Santa Barbara, CA, USA (Aug 17–21, 1997). <https://doi.org/10.1007/BFb0052231>
20. Hoffstein, J., Howgrave-Graham, N., Pipher, J., Silverman, J.H., Whyte, W.: NTRUSign: Digital signatures using the NTRU lattice. In: Joye, M. (ed.) CT-RSA 2003. LNCS, vol. 2612, pp. 122–140. Springer, Heidelberg, Germany, San Francisco, CA, USA (Apr 13–17, 2003). https://doi.org/10.1007/3-540-36563-X_9
21. Hoffstein, J., Howgrave-Graham, N., Pipher, J., Whyte, W.: Practical lattice-based cryptography: NTRUEncrypt and NTRUSign. pp. 349–390. ISC, Springer, Heidelberg, Germany (2010). <https://doi.org/10.1007/978-3-642-02295-1>
22. Hu, Y., Wang, B., He, W.: NTRUSign with a new perturbation. IEEE Transactions on Information Theory **54**(7), 3216–3221 (2008). <https://doi.org/10.1109/TIT.2008.924662>
23. Lin, X., Suzuki, M., Zhang, S., Espitau, T., Yu, Y., Tibouchi, M., Abe, M.: Cryptanalysis of the Peregrine lattice-based signature scheme. IACR Cryptology ePrint Archive p. 1628 (2023), <https://eprint.iacr.org/2023/1628>
24. MATZOV: Report on the security of LWE: Improved dual lattice attack (2022), <https://doi.org/10.5281/zenodo.6412487>
25. National Institute of Standards and Technology: NIST post-quantum cryptography standardization process (2016), <https://csrc.nist.gov/Projects/post-quantum-cryptography/round-4-submissions>
26. National Institute of Standards and Technology: NIST post-quantum cryptography standardization process: Additional signatures (2023), <https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures>
27. Nguyen, P.Q., Regev, O.: Learning a parallelepiped: Cryptanalysis of GGH and NTRU signatures. In: Vaudenay, S. (ed.) EUROCRYPT 2006. LNCS, vol. 4004, pp. 271–288. Springer, Heidelberg, Germany, St. Petersburg, Russia (May 28 – Jun 1, 2006). https://doi.org/10.1007/11761679_17

28. Plantard, T., Sipasseuth, A., Dumondelle, C., Susilo, W.: DRS. Tech. rep., National Institute of Standards and Technology (2017), available at <https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-1-submissions>
29. Postlethwaite, E.W., Virdia, F.: On the success probability of solving unique SVP via BKZ. In: Garay, J. (ed.) PKC 2021, Part I. LNCS, vol. 12710, pp. 68–98. Springer, Heidelberg, Germany, Virtual Event (May 10–13, 2021). https://doi.org/10.1007/978-3-030-75245-3_4
30. Prest, T.: A key-recovery attack against Mitaka in the t -probing model. In: Boldyreva, A., Kolesnikov, V. (eds.) PKC 2023, Part I. LNCS, vol. 13940, pp. 205–220. Springer, Heidelberg, Germany, Atlanta, GA, USA (May 7–10, 2023). https://doi.org/10.1007/978-3-031-31368-4_8
31. Prest, T., Fouque, P.A., Hoffstein, J., Kirchner, P., Lyubashevsky, V., Pornin, T., Ricosset, T., Seiler, G., Whyte, W., Zhang, Z.: FALCON. Tech. rep., National Institute of Standards and Technology (2022), available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>
32. Ritterhoff, S., Maringer, G., Bitzer, S., Weger, V., Karl, P., Schamberger, T., Schupp, J., Wachter-Zeh, A.: FuLeeca. Tech. rep., National Institute of Standards and Technology (2023), <https://csrc.nist.gov/csrc/media/Projects/pqc-dig-sig/documents/round-1/spec-files/FuLeeca-spec-web.pdf>
33. Ritterhoff, S., Maringer, G., Bitzer, S., Weger, V., Karl, P., Schamberger, T., Schupp, J., Wachter-Zeh, A.: FuLeeca: A Lee-based signature scheme. In: Esser, A., Santini, P. (eds.) Code-Based Cryptography - 11th International Workshop, CBCrypto 2023, Lyon, France, April 22-23, 2023, Revised Selected Papers. Lecture Notes in Computer Science, vol. 14311, pp. 56–83. Springer (2023). https://doi.org/10.1007/978-3-031-46495-9_4
34. Schnorr, C., Euchner, M.: Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Mathematical programming* **66**, 181–199 (1994). <https://doi.org/10.1007/BF01581144>
35. Seo, E.Y., Kim, Y.S., Lee, J.W., No, J.S.: Peregrine: Toward fastest FALCON based on GPV framework. *Cryptology ePrint Archive*, Report 2022/1495 (2022), <https://eprint.iacr.org/2022/1495>
36. Sommer, N., Feder, M., Shalvi, O.: Finding the closest lattice point by iterative slicing. *SIAM Journal on Discrete Mathematics* **23**(2), 715–731 (2009). <https://doi.org/10.1137/060676362>
37. Yu, Y., Ducas, L.: Learning strikes again: The case of the DRS signature scheme. In: Peyrin, T., Galbraith, S. (eds.) ASIACRYPT 2018, Part II. LNCS, vol. 11273, pp. 525–543. Springer, Heidelberg, Germany, Brisbane, Queensland, Australia (Dec 2–6, 2018). https://doi.org/10.1007/978-3-030-03329-3_18