How Secure are NTRU+ Keys? Analysing the Best Module Attacks

Paola de Perthuis[1] and Filip Trenkić[2]

[1] Centrum Wiskunde & Informatica (CWI), Amsterdam, the Netherlands
[2] King's College London, United Kingdom

Abstract A recent analysis of module-BKZ [Ducas–Engleberts–Perthuis, Asiacrypt'25] predicts subexponential speedups for lattice reduction on module lattices over cyclotomic fields whose conductor has small odd prime factors. This threatens NTRU+, the South Korean standard for key encapsulation, and related schemes; however, the short vectors one seeks to find are short in the coefficient geometry, whereas module-BKZ operates in the canonical geometry.

We introduce *relative canonical* and *relative coefficient* embeddings, realising a module over a cyclotomic field K as a module over any subfield L, and characterise when the two agree up to scaling: exactly when the radicals of their respective conductors c and ℓ satisfy $\mathrm{rad}(c)/\mathrm{rad}(\ell) \in \{1, 2\}$. Furthermore we quantify the distortion incurred in general when passing between embeddings, and hence identify the optimal subfield for module-BKZ and BKZ to use.

Estimating the cost of the best attack with simulation-based models [Perthuis–Trenkić, PQCrypto'26], we find that the best module attack saves only 4–10 blocksizes, costing NTRU+ keys at most three bits of security in the Core-SVP model.

We further extend the analysis of overstretched NTRU regimes [Ducas–Woerden, Asiacrypt'19] to the module setting, and find the fatigue point is reached sooner than with unstructured counterparts, though NTRU+ parameterisations remain safe.

Keywords: Module Lattices · NTRU · Coefficient Embedding · Key Recovery

1 Introduction

Starting in 2021, the South Korean National Security Research Institute (NSRI) and National Intelligence Service (NIS) launched a standardisation competition, which resulted in the selection of final signature and encryption schemes in January 2025 [Kpq25]: HAETAE, AIMer (for signatures); SMAUG-T and NTRU+ (public-key encryption and key encapsulation mechanisms).

At the end of 2025, an average-case analysis of the module-BKZ lattice-reduction algorithm [DEP25] suggested that subexponential speedups could be gained when looking for short elements in some module lattices, and in particular cyclotomic number fields with conductors divisible by small odd primes seemed

the most vulnerable. NTRU+ was directly concerned, as it instantiates the lattice problem NTRU [HPS98] with module lattices over such number fields.

In this article, we study how module-BKZ speedups could affect the security of NTRU+ secret keys, with general results applicable to other schemes for which these are sampled spherically in coefficient embeddings. Indeed, such draws over a non-power-of-two cyclotomic number field results in a distorted geometry, with security impacts left as an open question in [DEP25].

Moreover, a module-lattice over a number field K may be embedded in multiple ways, not only into the rationals, but also any other subfield of K. This will result in various performances of both structured and unstructured attacks, and we provide theoretical results on optimal embedding choices in both cases.

We then proceed to simulations providing precise heuristic estimations of the security of NTRU+ against key-recovery primal attacks. These show that the loss is only of a few bits, less than what could have been expected from previous articles analysing module-BKZ over non-power-of-two cyclotomics [PT26, DEP25], because of the distortion induced by coefficient sampling.

As a positive result, we extend to module-BKZ the analysis of overstretched regimes [ABD16, KF17, DW21, KK24], in which the integer modulus parameter q used in NTRU becomes so large that reduction algorithms detect a dense submodule faster than expected. Our heuristic analysis, which is both asymptotic and concrete simulation-based, shows that this dangerous regime is reached faster when making use of the module structure in NTRU+ schemes.

1.1 Related Literature

Attacks Using Module-Lattice Structures. In the last decade, a line of research explored accelerations of the retrieval of short elements in rank-one modules without exhibiting any impact on higher-rank problems, such as [CDPR16, CDW17, BEF+17]. Then in 2019, for instances of higher rank, Lee, Pellet-Mary, Stehlé and Wallet [LPSW19] proposed the first algorithm to perform lattice reduction generalising the LLL algorithm [LLL82] to modules over any number field (even when it is non-Euclidean), for which enough precomputations would have been performed to provide an approximate Closest Vector Problem oracle in a log-unit lattice corresponding to the field. This algorithm was later fully dequantised by Micheli and Micciancio [DM22]. The authors of [LPSW19] also proposed worst-case reductions between instances of different ranks, later generalised by Mukherjee and Stephens-Davidowitz [MS20].

A little later in 2019, Kirchner, Espitau and Fouque [KEF19] proposed heuristic variants of the LLL algorithm that make use of module structures recursing on towers of subfields, to in some cases improve on its complexity, at the cost of a slightly worse output basis quality.

Then, in 2024, Karenin and Kirshanova [KK24] developed an implementation of module-LLL for power-of-two conductor cyclotomic number fields, for which they heuristically studied the overstretched regime of NTRU instances, corresponding to rank-two modules with a secret dense rank-one submodule. Their extension of the study of such parameters to this type of structured lattices

concluded to a negative result, for which they did not find accelerations with respect to the unstructured case.

Overstretching. The overstretched regime of NTRU was first exhibited in 2016 by Albrecht, Bai and Ducas [ABD16], which exploit the algebraic structure of NTRU lattices to recover secret keys in subexponential or even polynomial time when the modulus q is sufficiently large. Shortly after, Kirchner and Fouque [KF17] showed that this algebraic structure is not needed: lattice reduction alone already behaves much better than expected in this regime, because the NTRU lattice contains an unusually dense sublattice of large dimension, generated by the secret key. Using the Pataki–Tural lemma [PT08], they concluded that moduli as small as $q = n^{2.783+o(1)}$ (where n is the rank of the sampled lattice over the rationals) already lie in the overstretched regime for ternary secrets. However, their impossibility argument only gives an asymptotic upper bound on this *fatigue point*, and does not explain how lattice reduction actually finds the dense sublattice.

In 2021, Ducas and van Woerden [DW21] identified the mechanism, by generalising the Pataki–Tural lemma to bound the volume of the intersection of the dense sublattice with the sublattices spanned by the first basis vectors. This reduced the asymptotic fatigue point down to $q = n^{2.484+o(1)}$. Furthermore, they developed a model to concretely estimate the fatigue point, and experimentally validated its accuracy. Finally, as mentioned above, Karenin and Kirshanova [KK24] proved an algebraic analogue of the Pataki–Tural lemma for modules over any number field, and used it to study the corresponding dense submodule discovery event for module-LLL of power-of-two cyclotomics.

NTRU+ Schemes. A few years ago, Kim and Parks [KP23] proposed a Key Encapsulation Mechanism (KEM) sampling rank-two modules over cyclotomic number fields with conductors of radical equal to six, instead of two as in most other lattice-based schemes. This scheme was later standardised at the beginning of 2025 by the South-Korean standardisation organisation, KpqC [KP25]. The first article was followed up by a Public-Key Encryption (PKE) construction [KP24], and the use of such cyclotomic number fields was extended to signatures [HWK+25]. In all of these publications, the security estimates were made without taking into account a possible discrepancy between unstructured and structured attacks; but [DEP25, PT26] suggest there should be subexponential accelerations of module-BKZ when reducing modules over such cyclotomic number fields.

1.2 Technical Overview

Considering Different Embeddings. Previous works [DEP25, PT26] analysed the performance of mBKZ with respect to BKZ on canonical embeddings of modules, and [PT26] considered short vectors that would be spherically sampled in this embedding. However in practice, constructions use vectors that are sampled spherically in a coefficient embedding. Moreover, module-BKZ may be run on

any subfield L of the field on which the original module is defined, resulting in several possible embedding choices from K to L. We theoretically compare the performance of choosing a natural *relative coefficient* embedding to an embedding that preserves the trace norm.

Distortion Induced by Embeddings. More particularly, we show that if K is a cyclotomic number field, then relative coefficient embeddings into subfields L will preserve the trace norm up to a scaling factor exactly when L is cyclotomic and all the odd primes in the conductor of K divide the conductor of L. We will call such subfields *conductor-odd-radical-preserving* or *distorsion-free*.

The performance of primal attacks increases when the length of the vector to be recovered diminishes. We show that when using a relative coefficient embedding of the module, this length scales in more than $\sqrt{d_L}$, where d_L is the degree of the subfield L used to embed, with equality for distortion-free and degree-one subfields, while the covolume of the lattice obtained also evolves with the discriminant of L.

Asymptotic Comparison of Primal Attack Strategies. We further analyse with usual heuristics the consequence of these changes in the performance of structured (using mBKZ) or unstructured (BKZ) primal attacks lead over these subfields using relative coefficient or canonical embeddings. We show that asymptotically the best module attacks will use cyclotomic subfields with conductor the odd radical of the conductor of the original field, and that using non-cyclotomic subfields would not induce better performances.

For unstructured BKZ, we show that with short vectors sampled spherically in a coefficient embedding, running BKZ on a canonical embedding of a relative coefficient embedding onto a cyclotomic subfield will be optimal if the subfield has a power-of-two conductor, which is in particular the case of $\mathbb{Q}$.

Application to NTRU+. Having identified the optimal subfield to choose, we employ fine-grained simulation-based security estimates (adapting [PT26]) and confirm that the geometry mistmatch blunts the improved slope of module-BKZ, shaving only a few bits of security in each parameterisation.

Overstretching. We further adapt the simulation-based estimates for overstretching [DW21] to the module setting. We borrow techniques from a previous generalisation of overstretching to module lattices [KK24], although the module-lattice reduction algorithm there considers an algebraic-SVP oracle, whereas module-BKZ works with a euclidean oracle. We make extensive use of results about filtrations and projections of module lattices proven in [MS20], and use this to provide a new proof of the module Pataki–Tural Lemma and a useful generalisation.

While the difference in fatigue point is invisible at the level of first-order asymptotics, we concretely determine a decrease when running module-BKZ on NTRU+ modules. Nevertheless, all NTRU+ parameterisations remain safe.

2 Preliminaries

Notations. $\mathbb{N}, \mathbb{Z}, \mathbb{Q}, \mathbb{R}$ and $\mathbb{C}$ will denote natural and relative integers, rational, real and complex numbers respectively. For n a natural integer, $[n]$ will refer to the set $\{1, 2, \ldots, n\}$, and $\mathcal{P}_n$ to the set of primes in $\mathbb{N}$ dividing n. $\subset$ will denote an inclusion (or equality) relation. For a ring R, $R^\times$ will denote its set of invertible elements for the multiplication operation. Given a matrix M, M^T will denote its transpose and M^* its conjugate transpose. Unless specified otherwise, vectors will be vertical.

Landau Asymptotic Notations. For functions f and g defined over an infinite subset S of $\mathbb{R}$ containing arbitrarily large positive values, with range admitting an absolute value $|\cdot|$, one writes

- $f(n) = o(g(n))$ if for every $\varepsilon > 0$, there exists a real n_0 such that for each n in S bigger than n_0, $|f(n)| \leq \varepsilon \cdot |g(n)|$;
- $f(n) \sim g(n)$ if for every $\varepsilon > 0$, there exists a real n_0 such that for each n in S bigger than n_0, $|f(n) - g(n)| \leq \varepsilon \cdot |g(n)|$;
- $f(n) = \Theta(g(n))$ if there exists $k, k' > 0$ and a real n_0 such that for each n in S bigger than n_0, $k \cdot |g(n)| < |f(n)| < k' \cdot |g(n)|$;
- $f(n) = O(g(n))$ if there exists $k > 0$ and a real n_0 such that for each n in S bigger than n_0, $|f(n)| \leq k \cdot |g(n)|$;
- $f(n) = \Omega(g(n))$ if $f(n) \neq o(g(n))$, *i.e.* there exist $k > 0$ such that for every real n_0, there exists n in S bigger than n_0, such that $|f(n)| \geq k \cdot |g(n)|$.

2.1 General Algebraic Background

Number Fields. We will define a number field as $K = \mathbb{Q}[X]/(f(X))$ for f a monic irreducible polynomial of $\mathbb{Z}[X]$, and call the degree of K the degree of f, corresponding to $[K : \mathbb{Q}]$. The c-th cyclotomic number field corresponds to the degree $\varphi(c)$ quotient polynomial $f = \Phi_c$, where φ denotes Euler's totient function and Φ_c the c-th cyclotomic polynomial; such a field is isomorphic to $\mathbb{Q}(\zeta_c)$ (that we will also call a cyclotomic number field), for ζ_c a primitive c-th root of unity.As discussed hereafter, for odd c greater than one, the relation $\Phi_{2c}(X) = \Phi_c(-X)$ defines an isomorphism between corresponding number fields. c is called the conductor of K. In general, for any number field K, the conductor is the smallest c such that K is included in the c-th cyclotomic field (existence is induced by the Kronecker-Weber theorem). For an integer n in $\mathbb{N}$ with prime decomposition $\prod_p p^{v_p(n)}$, the radical of n is $\mathrm{rad}(n) = \prod_{v_p(n)\neq 0} p$. In this work, we will call the *odd radical* $\prod_{p\neq 2 | v_p(n)\neq 0} p$.

A polynomial f defining a number field K admits $d_\mathbb{R}$ real roots and $2d_\mathbb{C}$ non-real roots; $(d_\mathbb{R}, d_\mathbb{C})$ is called the signature of K. We write μ_K for the set of roots of unity in K. A number field $K = \mathbb{Q}[X]/(f(X))$ is *totally real* iff f's roots are all real. K is a *CM* number field if none of the roots of f are in $\mathbb{R}$ and it contains a totally real subfield L such that K/L is a quadratic extension; one

then writes that K admits complex conjugation. Cyclotomic number fields of conductor c strictly greater than two are CM: indeed, primitive c-th roots of unity are not real, and, for ζ_c one of them, $[\mathbb{Q}(\zeta_c) : \mathbb{Q}(\zeta_c + \zeta_c^{-1})] = 2$, because ζ_c is a root of $Z^2 - (\zeta_c + \zeta_c^{-1})Z + 1 \in \mathbb{Q}(\zeta_c + \zeta_c^{-1})[Z]$.

The ring of integers $\mathcal{O}_K$ of a number field K denotes the set of algebraic integers of K, *i.e.* elements x in K such that there exists p in $\mathbb{Z}[X]$ for which $p(x) = 0$. If ζ in $\mathbb{C}$ is a root of unity, then $\mathcal{O}_{\mathbb{Q}(\zeta)} = \mathbb{Z}[\zeta]$ (see [Was97, Theorem 2.6]). This induces the following corollary.

Corollary 1. *If $K = \mathbb{Q}(\zeta)$ is a cyclotomic number field, and L one of its subfields, with ring of integers, then $\mathcal{O}_L[\zeta] = \mathcal{O}_K$.*

Proof. $\mathbb{Z}$ is a subset of $\mathcal{O}_L$, so $\mathbb{Z}[\zeta] = \mathcal{O}_K \subset \mathcal{O}_L[\zeta]$. Moreover $\mathcal{O}_L \subset \mathcal{O}_K$, so as $\mathcal{O}_K$ is a ring, $\mathcal{O}_L[\zeta] \subset \mathcal{O}_K$.

If K is a number field and L one of its subfields, there exists a monic irreducible degree-$[K : L]$ polynomial h in $L[Z]$ such that $K = L[Z]/(h(Z))$, that we will call the minimal polynomial of K on L.

Cyclotomic Fields of Odd Conductor. Denoting by Φ_n the n-th cyclotomic polynomial, for an odd natural c greater than one, $\Phi_c(X) = \Phi_{2k}(-X)$, and both are isomorphic to $\mathbb{Q}(\zeta_c)$. Generally, to define fields uniquely, a convention that no conductor is equal to two modulo four is adopted; however, this choice will have consequences in the geometry of certain non-canonical embeddings and will hence be discussed in further detail in the article.

Relative Trace. For K a number field and L one of its subfields, denoting $\overline{L}$ its algebraic closure, the trace $\mathrm{Tr}_{K/L}$ of K over L is defined as

$$\begin{aligned} \mathrm{Tr}_{K/L} \ : \ & K \to L \\ & a \mapsto \sum_{\sigma \in \mathrm{Hom}_L(K, \overline{L})} \sigma(a), \end{aligned}$$

where the destination space is L because it is the field fixed by $\mathrm{Aut}(\overline{L}/L)$. Alternatively, multiplication by a in K defines an endomorphism that may be represented by a matrix in a basis representing K in the vector space $L^{[K:L]}$, and $\mathrm{Tr}_{K/L}(a)$ is then the trace of such a matrix, invariant by basis change. Remark that for a tower of number fields $M \subset L \subset K$, and a in K,

$$\begin{aligned} \mathrm{Tr}_{L/M} \circ \mathrm{Tr}_{K/L}(a) &= \sum_{\sigma \in \mathrm{Hom}_M(L, \overline{M})} \sigma \left(\sum_{\tau \in \mathrm{Hom}_L(K, \overline{L})} \tau(a) \right) \\ &= \sum_{\substack{\sigma \in \mathrm{Hom}_M(L, \overline{M}) \\ \tau \in \mathrm{Hom}_L(K, \overline{L})}} \sigma \circ \tau(a) = \mathrm{Tr}_{K/M}(a). \end{aligned}$$

Discriminants and Algebraic Norm. For K a number field and L one of its subfields, the *relative discriminant* $\Delta_{K/L}$ ($\Delta_{\mathcal{O}_K/\mathcal{O}_L}$ would be more precise, but we use this notation for conciseness) is the ideal of $\mathcal{O}_L$ generated by the discriminants of all the bases of K/L contained in $\mathcal{O}_K$ [Neu99, III, Definition 2.8]. If $\mathcal{O}_K$ is a free $\mathcal{O}_L$-module of basis $(b_1, \ldots, b_{[K:L]}) \in K^{[K:L]}$ (*cf.* Section 2.2), then

$$\Delta_{K/L} = \det\left((\mathrm{Tr}_{K/L}(b_i b_j))_{i,j}\right).$$

Neukirch [Neu99, III, Theorem 2.9] explains that

$$\Delta_{K/L} = \mathrm{N}_{K/L}\left(\left\{x \in K \;\middle|\; \mathrm{Tr}_{K/L}\left(x\mathcal{O}_K\right) \subset \mathcal{O}_L\right\}\right)^{-1},$$

where the *algebraic norm* $\mathrm{N}_{K/L}$ is defined as

$$\begin{aligned} \mathrm{N}_{K/L} : K &\to L \\ x &\mapsto \prod_{\sigma \in \mathrm{Hom}_L(K, \overline{L})} \sigma(x), \end{aligned}$$

and $\mathrm{N}_{K/L}(x)$ also corresponds to the determinant of the endomorphism $a \mapsto xa$ of K as an L-vector space (see [Neu99, III, 2.5, 2.6]). In particular, for a tower of field extensions $M \subset L \subset K$, $\mathrm{N}_{K/M} = \mathrm{N}_{L/M} \circ \mathrm{N}_{K/L}$. When the field is clear from context, $\mathrm{N}_{K/\mathbb{Q}}$ will be denoted as N, which we will extend to $K_{\mathbb{R}} = K \otimes_{\mathbb{Q}} \mathbb{R}$ by tensorisation with $\mathbb{R}$. Similarly, Δ_K will be the shorthand for $\Delta_{K/\mathbb{Q}}$. One has $\Delta_{K/M} = \Delta_{L/M}^{[K:L]} \mathrm{N}_{L/M}(\Delta_{K/L})$ [Neu99, III.2.10].

If K is a cyclotomic number field of degree d, signature $(d_{\mathbb{R}}, d_{\mathbb{C}})$ and conductor c, then [Was97, 2.7]

$$\Delta_K = (-1)^{d_{\mathbb{C}}} \frac{c^d}{\prod_{p \in \mathcal{P}_c} p^{\frac{d}{p-1}}}.$$

Dirichlet Characters. For G a locally compact abelian group, a character of G is a continuous morphism of G onto the unit circle of $\mathbb{C}$ with the multiplication operation. The dual of G, denoted $\widehat{G}$, is the group of its characters, using pointwise multiplication. The Pontryagin theorem [Pon34] states that

$$G \to \widehat{\widehat{G}}, \quad g \mapsto (\chi \mapsto \chi(g))$$

is an isomorphism of topological groups.

If G is a finite abelian group, then $\widehat{G}$ is isomorphic to G [Was97, 3.1]. For H a subgroup of G, one defines the orthogonal of H as

$$H^{\perp} = \left\{\chi \in \widehat{G} \mid \forall h \in H, \chi(h) = 1\right\},$$

for which $H^{\perp} \simeq \widehat{G/H}$ and $\widehat{G}/H^{\perp} \simeq \widehat{H}$ [Was97, 3.3].

A *Dirichlet character* modulo a natural integer c is a character χ of the group $(\mathbb{Z}/c\mathbb{Z})^\times$. If c divides n, χ induces a homomorphism from $(\mathbb{Z}/n\mathbb{Z})^\times$ to $\mathbb{C}^\times$ by composition with the natural map from $(\mathbb{Z}/n\mathbb{Z})^\times$ to $(\mathbb{Z}/c\mathbb{Z})^\times$, allowing to see χ a character modulo n. The conductor of χ is the smallest integer f_χ such that χ is a character modulo f_χ. A character defined modulo its conductor is called primitive. By convention, unless stated otherwise, we will use primitive characters.

If K is a cyclotomic number field of conductor c, the Galois group $\mathrm{Gal}(K/\mathbb{Q})$ is isomorphic to $(\mathbb{Z}/c\mathbb{Z})^\times$, hence by identification the Dirichlet characters modulo c are the characters of this Galois group. Every subfield L of K corresponds to a subgroup H of $\mathrm{Gal}(K/\mathbb{Q})$, and to the group of characters $H^\perp$. The conductor ℓ of L is smallest common multiple of conductors of the elements of $H^\perp$, and the conductor-discriminant formula [Was97, 3.11] provides

$$|\Delta_L| = \prod_{\chi \in H^\perp} f_\chi.$$

2.2 Module Lattices

Canonical Embedding. For a number field K of signature $(d_\mathbb{R}, d_\mathbb{C})$, $K_\mathbb{R} = K \otimes_\mathbb{Q} \mathbb{R} \simeq \mathbb{R}^{d_\mathbb{R}} \times \mathbb{C}^{d_\mathbb{C}}$. The set of embeddings from K to $\mathbb{C}$ which fix $\mathbb{Q}$, written $\mathrm{Hom}_\mathbb{Q}(K, \mathbb{C})$, contains d elements. The map

$$\begin{aligned} \sigma_\mathbb{Q} : K &\to \mathbb{C}^d \\ x &\mapsto (\sigma(x))_{\sigma \in \mathrm{Hom}_\mathbb{Q}(K,\mathbb{C})} \end{aligned}$$

defines an isomorphism from K to $\sigma_\mathbb{Q}(K)$ which is the *canonical embedding*, and extends naturally to an isomorphism from $K_\mathbb{R}$ to $\sigma_\mathbb{Q}(K) \otimes_\mathbb{Q} \mathbb{R}$. For any $z \in K_\mathbb{R}$ define the complex conjugate $\overline{z} = \sigma_\mathbb{Q}^{-1}(\overline{\sigma_\mathbb{Q}(z)})$, and square root $z^{1/2} = \sigma_\mathbb{Q}^{-1}(\sigma_\mathbb{Q}(z)^{1/2})$, where $\sigma_\mathbb{Q}$ has been extended to $K_\mathbb{R}$ and the conjugation and square root operations are applied componentwise.

The field trace $\mathrm{Tr} : K_\mathbb{R} \to \mathbb{R}, x \mapsto \sum_{j=1}^d \sigma_j(x)$ induces a Hermitian form $\langle x, y \rangle_\mathrm{Tr} = \mathrm{Tr}(x\bar{y})$, which coincides with $\langle \sigma_\mathbb{Q}(x), \sigma_\mathbb{Q}(y) \rangle_\mathbb{C}$. This extends to $K_\mathbb{R}^r$ by $\langle \mathbf{x}, \mathbf{y} \rangle_\mathrm{Tr} = \sum_{i=1}^r \langle x_i, y_i \rangle_\mathrm{Tr}$. On the other hand, we also define $\langle \mathbf{x}, \mathbf{y} \rangle_K = \sum_{i=1}^r x_i \bar{y}_i$. We write $\|\mathbf{x}\| = \langle \mathbf{x}, \mathbf{x} \rangle_\mathrm{Tr}^{1/2}$ and $\|\mathbf{x}\|_K = \langle \mathbf{x}, \mathbf{x} \rangle_K^{1/2}$. The algebraic norm satisfies, by the arithmetic-geometric inequality, $\sqrt{d}\mathrm{N}(x)^{1/d} \leq \|x\|$. If K is a CM or totally real field, we will use for $\mathbf{x}$ in $K_\mathbb{R}^m$ the shorthand $\mathrm{N}(\mathbf{x}) = \mathrm{N}(\langle \mathbf{x}, \mathbf{x} \rangle_K)^{1/2}$.

Modules. A *fractional ideal* of $\mathcal{O}_K$ is an $\mathcal{O}_K$-submodule $\mathfrak{a} \subset K$ for which there exists $x \in K \setminus \{0\}$ such that $x\mathfrak{a} \subseteq \mathcal{O}_K$. The algebraic norm is $\mathrm{N}(\mathfrak{a}) = [\mathcal{O}_K : x\mathfrak{a}]/|\mathrm{N}(x)|$ for any such x. A rank r $\mathcal{O}_K$-module is a set of the form $M = \sum_{i=1}^r \mathbf{b}_i \mathfrak{b}_i$ for nonzero fractional ideals $\mathfrak{b}_1, \ldots, \mathfrak{b}_r$ and $K_\mathbb{R}$-linearly independent vectors $\mathbf{b}_1, \ldots, \mathbf{b}_r \in K_\mathbb{R}^m$. M is said to be a *free* module if $\mathfrak{b}_1 = \cdots = \mathfrak{b}_r = \mathcal{O}_K$. The set $\mathfrak{B} = (\mathbf{b}_i, \mathfrak{b}_i)_{i=1}^r$ is a *pseudobasis* of M. We write $\mathbf{B}$ for the matrix with columns $\mathbf{b}_1, \ldots, \mathbf{b}_r$. The module has $K_\mathbb{R}$-determinant $\det_{K_\mathbb{R}}(M) = \det(\overline{\mathbf{B}}^T \mathbf{B}) \prod_{i=1}^r \mathfrak{b}_i$.

If we extend the canonical embedding $\sigma_{\mathbb{Q}} : K \to \mathbb{C}^d$ componentwise to elements of M, then $\sigma_{\mathbb{Q}}(M)$ is an rd-dimensional lattice in $\mathbb{C}^{rd} \simeq \mathbb{R}^{2rd}$, called a *module lattice*. This embedding preserves geometry, since the trace product on $K_{\mathbb{R}}^r$ coincides exactly with the Euclidean inner product on $\mathbb{C}^{rd}$,

$$\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^{r} \mathrm{Tr}_{K/\mathbb{Q}}(x_i \bar{y_i}) = \sum_{i=1}^{r} \sum_{\sigma \in \mathrm{Hom}_{\mathbb{Q}}(K,\mathbb{C})} \sigma(x_i)\overline{\sigma(y_i)} = \langle \sigma_{\mathbb{Q}}(x), \sigma_{\mathbb{Q}}(y) \rangle.$$

We will abuse notation and write M to mean both the $\mathcal{O}_K$-module and the module lattice. The $\mathbb{Q}$-determinant is the volume of the module lattice, $\det_{\mathbb{Q}}(M) = |\Delta_K|^{r/2} \mathrm{N}(\det_{K_{\mathbb{R}}}(M))$.

For a pseudobasis $\mathfrak{B} = (\mathbf{b}_i, \mathfrak{J}_i)_{i=1}^r$ of a module lattice $\mathcal{M}$ we define π_i as the $K_{\mathbb{R}}$-linear projection away from the span of $\{\mathbf{b}_1, \ldots, \mathbf{b}_{i-1}\}$. For $1 \le i < j \le r$ we define the *projected subbasis* $\mathfrak{B}_{[i:j]} = (\pi_i(\mathbf{b}_k), \mathfrak{J}_k)_{k=i}^j$ and the corresponding *projected submodule* $\mathcal{M}_{[i:j]} = \pi_i(\mathbf{b}_i)\mathfrak{J}_i + \ldots + \pi_i(\mathbf{b}_j)\mathfrak{J}_j$.

Definition 1 (K-profile). *The pseudobasis $\mathfrak{B} = (\mathbf{b}_i, \mathfrak{J}_i)_{i=1}^r$ has K-profile $(\ell_i^K)_{i=1}^r$ where $\ell_i^K = \ln \det_{\mathbb{Q}}(\mathbf{b}_i^* \cdot \mathfrak{J}_i)$.*

Definition 2 (Module Gaussian Heuristic [GSVV24, DEP25]). *Let M be a rank r module lattice of unit volume over $K \simeq \mathbb{Q}(\omega_c)$. The Gaussian heuristic for M approximates the first minimum as*

$$\mathrm{gh}_K(r) = \mathrm{gh}_{\mathbb{Q}}(rd) \cdot \left(\frac{|\mu_K|}{2}\right)^{1/rd}.$$

We write $\mathrm{lgh}_K(\cdot) = \ln \mathrm{gh}_K(\cdot)$. By appropriate scaling we define the module Gaussian heuristic for arbitrary M as $\mathrm{gh}_K(M) = \mathrm{gh}_K(r) \cdot \det_{\mathbb{Q}}(M)^{1/rd}$.

Filtrations. We have seen module lattices can be described in terms of pseudobasis, but [MS20] show that a natural alternative viewpoint is the language of *filtrations*.

Definition 3 (Filtration). *For M an $\mathcal{O}_K$-module of rank r, a filtration is a chain of module lattices $M_1 \subset \cdots \subset M_k = M$ such that*

- *each M_i is a primitive submodule of M, or explicitly $M_i = M \cap \mathrm{span}_{K_{\mathbb{R}}}(M_i)$,*
- *the rank increments along the chain, $\mathrm{rank}_{K_{\mathbb{R}}}(M_i) = i$,*
- *each linear projection $\widetilde{M_i} = \pi_i(M_i)$ orthogonally (for the inner-product over $K_{\mathbb{R}}^r$) to previous elements in the filtration is an $\mathcal{O}_K$-module of rank 1.*

In this work we use the term *filtration tail* to refer to a suffix of any filtration

Lemma 1 ([MS20, Fact 3.2]). *For any filtration tail $M_s \subset \ldots \subset M_k = M$ for an $\mathcal{O}_K$-module M, we have $\det_{\mathbb{Q}}(M) = \det_{\mathbb{Q}}(M_s) \prod_{i=s+1}^{k} \det_{\mathbb{Q}}(\widetilde{M_i})$.*

Coefficient Embedding. We have an alternative way to view a module lattice as a rational lattice, commonly used in cryptography, which does not preserve geometry. Write $K = \mathbb{Q}[X]/(f(X))$, then viewing any $a \in K$ as a polynomial $a = \sum_{i=0}^{d-1} a_i X^i$ with $a_i \in \mathbb{Q}$ gives an additive isomorphism

$$\begin{aligned} \kappa_{\mathbb{Q}} : K &\to \mathbb{Q}^d \\ \textstyle\sum_{i=1}^{d} a_i X^i &\mapsto (a_0, \ \cdots, a_{d-1})^T \end{aligned}$$

which is the *coefficient embedding* and extends naturally to $K_{\mathbb{R}}^r$. While $\mathrm{vec}(M) \subset \mathbb{R}^{rd}$ is indeed a lattice, the geometry is not preserved. We further define the ring homomorphism

$$\begin{aligned} \mathrm{rot} : K &\to \mathbb{Q}^{d \times d} \\ a &\mapsto (\kappa_{\mathbb{Q}}(a) \mid \kappa_{\mathbb{Q}}(aX) \mid \cdots \mid \kappa_{\mathbb{Q}}(aX^{d-1})) \end{aligned}$$

which sends each $a \in K$ to a matrix corresponding to the action of multiplication by a, in particular $\kappa_{\mathbb{Q}}(xy) = \mathrm{rot}(x)\kappa_{\mathbb{Q}}(y)$ and $\mathrm{rot}(xy) = \mathrm{rot}(x)\mathrm{rot}(y)$. If M admits an $\mathcal{O}_K$-basis $\mathbf{B} \in \mathcal{O}_K^{r \times r}$, then the block matrix $\mathrm{rot}(\mathbf{B}) \in \mathbb{R}^{rd \times rd}$ is a basis for the lattice $\kappa_{\mathbb{Q}}(M)$.

3 Relative Embeddings into Subfields

We now generalise embeddings into $\mathbb{Q}$-lattices as by considering a subfield $L \subset K$ and defining *relative canonical* and *relative coefficient* embeddings allowing us to view rank-ρ_K $\mathcal{O}_K$-modules as rank-$[K:L]\rho_K$ $\mathcal{O}_L$-modules.

Definition 4 (Trace-Preserving Embedding). *We will call a relative embedding σ of K a CM or totally real field into $L_{\mathbb{R}}^{[K:L]}$, for L a subfield of K,* trace-preserving *if, for any x, y in K, and $\sigma(x) = (\sigma_1(x), \ldots, \sigma_{[K:L]}(x))$, one always obtains*

$$\mathrm{Tr}_{K/\mathbb{Q}}(x\overline{y}) = \mathrm{Tr}_{L/\mathbb{Q}}\left(\sum_i \sigma_i(x)\overline{\sigma_i(y)}\right).$$

Definition 5 (Relative Canonical Embedding). *For L a subfield of $K = L[Z]/(h(Z))$ a CM or totally real field, we define the relative embedding σ_L as a map*

$$\begin{aligned} \sigma_L : K &\to L_{\mathbb{R}}^{[K:L]} \\ x &\mapsto \phi\left((\tau(x))_{\tau \in \mathrm{Hom}_L(K, \overline{L})}\right) \end{aligned}$$

where ϕ is an isometry from $(\tau(K))_{\tau \in \mathrm{Hom}_L(K,\overline{L})}$ to $L_{\mathbb{R}}^{[K:L]}$. In this article, σ_L is instantiated in the following way: denoting T^*T the Cholesky decomposition of

the matrix $(\mathrm{Tr}_{K/L}(\overline{Z^{i-1}}Z^{j-1}))_{i,j\in[[K:L]]}$,

$$\begin{aligned}\sigma_L : \quad & K \to L_{\mathbb{R}}^{[K:L]} \\ & x = \sum_{i=1}^{[K:L]} x_i Z^{i-1} \mapsto T \begin{pmatrix} x_1 \\ \vdots \\ x_{[K:L]} \end{pmatrix}.\end{aligned}$$

By definition, σ_L is trace-preserving.

Proof (Consistency of Definition 5). Denoting $\mathbf{x} = (x_1, \ldots, x_{[K:L]})$, there is a Vandermonde matrix $V = \left(\tau(Z^{j-1})\right)_{\tau\in\mathrm{Hom}_L(K,\overline{L}),j\in[[K:L]]}$ in $\mathcal{G}\ell_{[K:L]}(\overline{L})$ such that $(\tau(x))_{\tau\in\mathrm{Hom}_L(K,\overline{L})} = V\mathbf{x}$. We need to prove the existence of an orthogonal matrix O in $\mathcal{O}_{[K:L]}(\overline{L}_{\mathbb{R}})$ such that $T = OV$, or equivalently, that $V^*V = T^*T$. By definition,

$$\begin{aligned}T^*T &= \left(\mathrm{Tr}_{K/L}(\overline{Z^{i-1}}Z^{j-1})\right)_{i,j\in[[K:L]]} \in M_{[K:L]\times[K:L]}(L) \\ &= \left(\sum_{\tau\in\mathrm{Hom}_L(K,\overline{L})} \tau(\overline{Z^{i-1}}) \cdot \tau(Z^{j-1})\right)_{i,j\in[[K:L]]} \\ &= V^*V.\end{aligned}$$

□

Definition 6 (Relative Coefficient Embedding). *The relative coefficient embedding from K to L, where $K = L[Z]/(h(L))$, is the map*

$$\begin{aligned}\kappa_L : \quad & K \to L^{[K:L]} \\ & \sum_{i=1}^{[K:L]} x_i Z^{i-1} \mapsto \left(x_1, \ldots, x_{[K:L]}\right)^T.\end{aligned}$$

Remark 1 (Cyclotomic Subfields of Odd Conductor). Alternatively, in the case where two divides c but not ℓ, as $L = \mathbb{Q}[Y]/(\Phi_\ell(Y)) \simeq L' = \mathbb{Q}[Y]/(\Phi_{2\ell}(Y))$ through $\Phi_{2\ell}(Y) = \Phi_\ell(-Y)$, one may with a change of variable instead consider the map $\kappa_{L'}$ as L'.

The map presented above has appeared in the literature under different names, such as the *descent procedure* of [KK24]. We note it specialises to the usual coefficient embedding for $L = \mathbb{Q}$.

4 Distortion between Embeddings when Sampling Spherically in the Coefficient Embedding

A distribution for a module vector over a cyclotomic number field K that is spherical in its coefficient embedding $\kappa_{\mathbb{Q}}$ does not in general induce a spherical

distribution in other embeddings. This section shows what the distribution of the induced Euclidean norms of the vector in various embeddings will behave, as well as the conditions under which relative embeddings are not distorted.

For completeness, even though it is not used in this article, but part of the module-lattice literature such as [DD12], a study of the canonical embeddings norms obtains from sampling in a cyclic embedding is available in Appendix A.1.

4.1 Distortion in the Canonical Embedding

Let $A_1, \ldots, A_{d_K}$ be independent variables sampled according to a centred distribution $\mathcal{D}$ over $\mathbb{Q}$ of variance σ, yielding the element $A(X) = \sum_i A_i X^{i-1}$ in $K \simeq \mathbb{Q}(\zeta_c)$ a cyclotomic number field of conductor c and degree d_K. For each index j of $[c]$ that is coprime with c, one sets

$$B_j = A(\zeta_c^j).$$

The random vector $B = (B_j)_j$ is defined with a Vandermonde matrix as

$$B = \begin{pmatrix} 1 & \omega_1 & \cdots & \omega_1^{d_K-1} \\ 1 & \omega_2 & \cdots & \omega_2^{d_K-1} \\ \vdots & & & \vdots \\ 1 & \omega_d & \cdots & \omega_{d_K}^{d_K-1} \end{pmatrix} \begin{pmatrix} A_1 \\ \vdots \\ A_{d_K} \end{pmatrix},$$

denoting $(\omega_i)_i$ the Galois conjugates of ζ_c. For j and k in $[c]$ coprime with c,

$$\begin{aligned} B_j \overline{B}_k &= \sum_{i,\ell=1}^{d} A_i \overline{A}_\ell \zeta_c^{j(i-1)} \overline{\zeta}_c^{k(\ell-1)} \\ &= \frac{1}{c} \left(\sum_i A_i^2 (\zeta_c^j \overline{\zeta}_c^k)^{i-1} + \sum_{i \neq \ell} A_i A_\ell \zeta_c^{j(i-1)} \zeta_c^{-k(\ell-1)} \right), \end{aligned}$$

hence by linearity of the expected value, independence of the A_i variables, and as $\zeta_c^j \overline{\zeta}_c^k$ is a c-th root of unity,

$$\mathrm{Cov}(B_j, B_k) = \begin{cases} \sigma^2 \sum_{i=1}^{d_K} (\zeta_c^{j-k})^{i-1} & \text{if } k \neq j \\ d\sigma^2 & \text{otherwise,} \end{cases}$$

while by linearity, the expected value of B is a null vector. For ζ_c^j a primitive c-th root of unity,

$$B_j \overline{B_j} = \sum_{k=1}^{d_K} A_k^2 + \sum_{\substack{k,\ell=1 \\ k \neq \ell}}^{d_K} A_k A_\ell \zeta_c^{j(k-\ell)},$$

thus the Euclidean norm of B is provided by

$$\|B\| = \sqrt{d_K \sum_{k=1}^{d_K} A_k^2 + \sum_{\substack{k,\ell=1 \\ k\neq\ell}}^{d_K} A_k A_\ell \sum_{\substack{j\in[c] \\ j\wedge c=1}} \zeta_j^{k-\ell}},$$

and in particular, the expected value of $\|B\|^2$ is $d_K^2\sigma^2$:

$$\mathbb{E}\left[\|\sigma_{\mathbb{Q}}(A)\|^2\right] = d_K^2\sigma^2.$$

4.2 Distortion in a Relative Coefficient Embedding

For $K = \mathbb{Q}[X]/(f(X)) \simeq \mathbb{Q}(\zeta_c)$ a cyclotomic number field of conductor c and $L = \mathbb{Q}[Y]/(g(Y))$ on of its subfields of degree d_L, such that $K \simeq L[Z]/(h(Z))$. Let $A_1, \ldots, A_{d_K}$ be independent variables sampled from a centred distribution $\mathcal{D}$ of standard deviation σ over $\mathbb{Q}$. There exists $B_{1,1}, \ldots, B_{1,d_L}, \ldots, B_{[K:L],d_L}$ random variables over $\mathbb{Q}$ such that, for each complex α that is a root of f, there exist respective roots β and γ of g and h such that $\sum_{i=1}^{d_K} A_i\alpha^{i-1} = \sum_{i,j} B_{i,j}\beta^{j-1}\gamma^{i-1}$, and there exists an injective map ϕ from roots of f to a pair of roots of g and h respectively such that $\phi(\alpha) = (\beta, \gamma)$ Solving d_K linear equations over $\mathbb{C}$ provides the $B_{i,j}$ from the A_k. In particular, there exists $M \in \mathcal{G}\ell_{d_K}(\mathbb{C})$ such that, denoting $B = (B_{1,1}, \ldots, B_{1,d_L}, \ldots, B_{[K:L],d_L}), A = (A_1, \ldots, A_{d_K})$, $B = MA$. Then, by linearity, the $B_{i,j}$ are also centred variables, and for each $(i,j) \in [[K:L]] \times [d_L]$, denoting $M_{i,j}$ the corresponding line of M, $\mathbb{E}[B_{i,j}^2] = \sigma^2\|M_{i,j}\|^2$. Then, denoting $\|\cdot\|_F$ the Frobenius norm of a matrix,

$$\|\kappa_L(A)\| = \sqrt{\sum_{i=1}^{[K:L]} \left(d_L \sum_{k=1}^{d_L} B_{i,k}^2 + \sum_{\substack{k,l=1 \\ k\neq l}}^{d_L} B_{i,l}B_{i,k} \sum_{\substack{j\in[\ell] \\ j\wedge\ell=1}} \zeta_\ell^{j(k-l)} \right)}$$

$$\mathbb{E}\left[\|\kappa_L(A)\|^2\right] = \sigma^2 d_L \sum_{i=1}^{[K:L]} \sum_{j=1}^{d_L} \|M_{i,j}\|^2 = \sigma^2 d_L \mathrm{Tr}(M^T M) = \sigma^2 d_L \|M\|_F$$

$$\mathbb{E}\left[\|\kappa_L(A)\|^2\right] = d_L\sigma^2 \sum_{i=1}^{d_K} \|\kappa_L(\zeta_c^{i-1})\|^2.$$

One may lower bound the Frobenius norm of M, provided by $\mathrm{Tr}(M^T M)$, using its determinant. In particular, $M = N^{-1}V$, where V is the Vandermonde matrix corresponding to K, with lines $(1, \omega_c, \ldots, \omega_c^{d_K-1})$, where the ω_c are the Galois conjugates of ζ_c, and N has rows of the form $(1, \lambda, \ldots, \lambda^{d_L-1}, \gamma, \lambda\gamma, \ldots, \lambda^{d_L-1}\gamma^{[K:L]-1})$, where for each line index corresponding to an ω_c in V, the corresponding line of N uses $(\lambda, \gamma) = \phi(\omega_c)$. The determinant of V corresponds to the covolume of $\mathcal{O}_K$, with $\det(V)^2 = \Delta_K$ and $\mathrm{vol}(\mathcal{O}_K) = \sqrt{|\Delta_K|} = |\det(V)|$, and the one of N to the covolume of $\mathcal{O}_L[\zeta_c]$, because just as the columns of V provide a $\mathbb{Z}$-basis

of $\mathcal{O}_K$ via its Minkowski embedding, those of N provide one of $\mathcal{O}_L[\zeta_c]$. One has $\det(M)^2 = [\mathcal{O}_K : \mathcal{O}_L[\zeta_c]]^2$. As K is cyclotomic, $\mathcal{O}_L[\zeta_c]$ is equal to $\mathcal{O}_K$. Hence, $\det(M)^2 = 1$, and $\|M\|_F \geq d_K$. By arithmetic-geometric inequality,

$$\mathbb{E}\left[\|\kappa_L(A)\|^2\right] \geq d_K d_L \sigma^2,$$

with equality iff M is orthogonal, corresponding to the conditions described in the following lemma.

Lemma 2 (Minimal-Norm Relative Coefficient Embedding for a Given Degree). *Let K be a cyclotomic number field of conductor c, L be one of its subfields, $A_1, \ldots, A_{d_K}$ be independent random variables samples from a centred distribution of standard deviation σ over $\mathbb{Q}$, and $A = \sum_i A_i X^{i-1}$ be in K. Then,*

$$\mathbb{E}[\|\kappa_L(A)\|^2] = d_K d_L \sigma^2$$

if and only if: (i) L is a cyclotomic subfield of K of degree greater than one and conductor ℓ such that $\mathrm{rad}(c)/\mathrm{rad}(\ell)$ is in $\{1; 2\}$, or (ii) L is of degree one.

Proof. See Appendix A.2.

Lemma 3 (Cyclotomic Subfields Preserving the Conductor Radical). *For K a cyclotomic field of conductor c and $L = \mathbb{Q}[Y]/(\Phi_\ell(Y))$ a cyclotomic subfield of K,*

(i) $\mathrm{N}_{L/\mathbb{Q}}(\Delta_{K/L}) = \frac{|\Delta_K|}{|\Delta_L|^{[K:L]}} = [K:L]^{\varphi(c)} \iff \frac{\mathrm{rad}(c)}{\mathrm{rad}(\ell)} \in \{1; 2\}$;

(ii) $\frac{c}{\ell} = \frac{\varphi(c)}{\varphi(\ell)} \iff \mathrm{rad}(\ell) = \mathrm{rad}(c) \iff K = L[Z]/(Z^{[K:L]} - Y)$;

(iii) $\frac{c}{\ell} = \frac{2\varphi(c)}{\varphi(\ell)} \iff \mathrm{rad}(c) = 2 \cdot \mathrm{rad}(\ell)$

Remark 2. When $\mathrm{rad}(c) = 2 \cdot \mathrm{rad}(\ell)$, K is equal to $L[Z]/(Z^{[K:L]} + Y)$, as may be observed from (ii) through a variable change from ℓ to 2ℓ and Y to $-Y$.

Proof. With the prime decompositions $c = \prod_{p \in \mathcal{P}_c} p^{e_p}, \ell = \prod_{p \in \mathcal{P}_\ell \subset \mathcal{P}_c} p^{f_p}$, $\frac{\varphi(c)}{\varphi(\ell)} = \frac{\prod_{p \in \mathcal{P}_c}(p-1)p^{e_p-1}}{\prod_{p \in \mathcal{P}_\ell}(p-1)p^{f_p-1}} = \frac{c}{\ell} \cdot \frac{\varphi(\mathrm{rad}(c)/\mathrm{rad}(\ell))}{\mathrm{rad}(c)/\mathrm{rad}(\ell)}$, which are equal to $\frac{c}{\ell}$ iff $\mathrm{rad}(c) = \mathrm{rad}(\ell)$ and to $\frac{c}{2\ell}$ iff $\mathrm{rad}(c) = 2 \cdot \mathrm{rad}(\ell)$. Moreover, using the expression of cyclotomic field discriminants,

$$\frac{\mathrm{N}_{L/\mathbb{Q}}(\Delta_{K/L})^{\frac{1}{\varphi(c)}}}{[K:L]} = \prod_{p \in \mathcal{P}_c \setminus \mathcal{P}_\ell} \frac{p^{1-\frac{1}{p-1}}}{p-1},$$

and for $x > 1$,

$$\frac{d\left(\left(1 - \frac{1}{x-1}\right)\ln x - \ln(x-1)\right)}{dx} = \frac{\ln x}{(x-1)^2} - \frac{1}{x(x-1)}$$

which is stricly positive for $x \geq 2$, and as $\lim_{x\to+\infty} \ln \frac{x^{1-\frac{1}{x-1}}}{x-1} = 0$, the derivative is negative on $[2;+\infty[$, and as $3^{1-1/2}/2 < 1$, $\mathrm{N}_{L/\mathbb{Q}}(\Delta_{L/K})$ is equal to $[K:L]^{\varphi(c)}$ iff $\frac{\mathrm{rad}(c)}{\mathrm{rad}(\ell)} \in \{1;2\}$.

If $\frac{c}{\ell} = \frac{d_K}{d_L}$, then $\zeta_c^{[K:L]} = \zeta_c^{c/\ell} = \zeta_\ell$, and ζ_c is hence a root of the degree-$[K:L]$ monic polynomial $Z^{[K:L]} - Y \in L[Z]$. Hence, the minimal polynomial of L on K divides $Z^{[K:L]} - Y$, and as it is monic of the same degree it is equal to $Z^{[K:L]} - Y$, which is irreducible. Finally, $K \simeq L(\zeta_c)$ is isomorphic to $L[Z]/(Z^{[K:L]} - L)$.

If $K = L[Z]/(Z^{[K:L]} - Y)$, as ζ_c is a root of $Z^{[K:L]} - Y$, $\zeta_c^{d_K/d_L}$ is a primitive ℓ-th root, which means that there exists k coprime with ℓ such that $\frac{d_K}{d_L c} = \frac{k}{\ell}$, hence such that $k \cdot \mathrm{rad}(c)/\mathrm{rad}(\ell) = \varphi(\mathrm{rad}(c)/\mathrm{rad}(\ell))$. As for any integer x, $\varphi(x) \leq x$, k is equal to one and $\mathrm{rad}(c)/\mathrm{rad}(\ell)$. □

Relative Embeddings Without Distortion There is no distortion between the relative canonical and coefficient embeddings if and only if L is an odd-radical preserving cyclotomic subfield of K.

Theorem 1 (Relative Coefficient Embeddings without Distortion). *For K a cyclotomic number field of conductor c and L a subfield of K of conductor ℓ, relative coefficient and canonical embeddings from K to L agree up to a constant scaling factor, if and only if L is cyclotomic and $\frac{\mathrm{rad}(c)}{\mathrm{rad}(\ell)} \in \{1;2\}$. In this case,*

$$\|\sigma_L\|^2 = [K:L] \cdot \|\kappa_L\|^2.$$

Remark that for L of degree one, one finds, as in usual literature, that the coefficient embedding $\kappa_{\mathbb{Q}} : K \to \mathbb{Q}^{[K:\mathbb{Q}]}$ is an scaled isometry for $K \simeq \mathbb{Q}(\omega_c)$ if and only if c is a power of two.

Proof. See Appendix A.3.

5 Asymptotic Comparison of Primal Attack Predictions

Section 5.1 proposes an asymptotic modelling of the minimal blocksize of (module) primal attacks against a uSVP instance provided by a module lattice over a field K for which the short vector $\mathbf{v}$ was spherically sampled in the coefficient embedding. This makes use of Gaussian, (module)-GSA heuristics, the beta model [PT26] for the norm of projections of $\mathbf{v}$, and parameter regimes in which the (module)-GSA is believed to hold (with a blocksize less than half of the total rank of the embedded lattice on which the attack is run).

Section 5.2 then analyses the consequences of the model when comparing different attack strategies, underlining the importance of relative discriminant terms between K and L, and Section 5.3 then concludes when K is cyclotomic, after proving bounds on the relevant discriminant terms, that the optimal module attack will use a relative coefficient (or canonical) embedding of the module into a cyclotomic subfield with the same conductor odd-radical as K, and the

optimal unstructured attack the coefficient embedding of the module. Finally, the asymptotic gain of the structured attack with respect to the classical one is slightly subexponential, with a $|\Delta_K|^{1/d_K}/d_K$ discriminant-term constant, denoting Δ_K and d_K the discriminant and degree of K respectively.

5.1 Asymptotic Modelling of Primal Attacks

An attacking algorithm $\mathcal{A}$ on a module-uSVP instance in the form of a rank-ρ_K module M over the number field K of degree d_K, which emcompasses $\sqrt{\delta_{\mathcal{A}}}\mathbf{v}\mathcal{O}_K$ where $\mathbf{v}$'s coordinates are taken from a distribution of elements of short norm, using, for L a degree-d_L subfield of K, mBKZ with blocksize $\beta_{\mathcal{A}}/d_L$ on an embedding of K into $L^{[K:L]}$, or BKZ with blocksize $\beta_{\mathcal{A}}$, succeeds with probability more than one half if, under the Gaussian and GSA heuristics and the beta-model for the length of projected vectors,

$$\begin{aligned}&\frac{1}{2}\mathbb{E}\left[\ln\left(\delta_{\mathcal{A}}\|\mathbf{v}\|^2\right)\right] + \frac{1}{2}\left(\psi\left(\frac{\beta_{\mathcal{A}}}{2}\right) - \psi\left(\frac{d_K\rho_K}{2}\right)\right)\\ &\quad\leq \mathrm{lgh}_L\left(\frac{\beta_{\mathcal{A}}}{d_L}\right) + \frac{d_K\rho_K - \beta_{\mathcal{A}}}{2}\cdot\mathbb{Q}\text{-slope}_L(\beta_{\mathcal{A}}) + \frac{\ln D_{\mathcal{A}}}{2d_K} + \frac{\ln Q}{2},\end{aligned} \tag{1}$$

where the volume over $\mathbb{Q}$ of the embedding of M on which the attack is run is denoted as $\left(\sqrt{D_{\mathcal{A}}^{1/d_K}Q}\right)^{\rho_K d_K}$, and the attacking-algorithm-dependent parameters $\delta_{\mathcal{A}}$ and $D_{\mathcal{A}}$ are described in Table 1. Remark that BKZ on $\sigma_{\mathbb{Q}}(M)$ is mBKZ on $\sigma_{\mathbb{Q}}(M)$ for the subfield $\mathbb{Q}$ of K.

Attack	$\delta_{\mathcal{A}}$	$D_{\mathcal{A}}$
BKZ on $\kappa_L(M)$	$\geq d_L$	$\lvert\Delta_L\rvert^{[K:L]}$
mBKZ on $\kappa_L(M)$	$\geq d_L$	$\lvert\Delta_L\rvert^{[K:L]}$
BKZ on $\sigma_{\mathbb{Q}}(M)$	d_K	$\lvert\Delta_K\rvert$
mBKZ on $\sigma_L(M)$	d_K	$\lvert\Delta_K\rvert$

Table 1: Attack-Strategy Parameters. The lower bounds d_L on $\delta_{\mathcal{A}}$ are attained when L fulfils (i) of Lemma 3, or is of degree one (*cf.* Lemma 2).

Heuristic 1 (Asymptotic Modelling of Primal Attacks). *For parameters fulfilling an equality in relation 1, when $\beta_{\mathcal{A}}$ tends to infinity we model $\mathbb{E}\left[\ln\|\mathbf{v}\|\right] = \frac{1}{2}\ln(\rho_K d_K) + \ln(\sigma) + O\left(\frac{1}{\rho_K d_K}\right)$, where σ is a strictly positive variable specified in the distribution of $\mathbf{v}$ that converges to a finite limit, L and $\ln(D_{\mathcal{A}})/\delta_{\mathcal{A}} - \ln(\delta_{\mathcal{A}})$ as constant, $\rho_K d_K > 2\beta_{\mathcal{A}}$, and $Q = o(\beta_{\mathcal{A}})$. These conditions imply that $Q = \Omega(\sigma^2\beta_{\mathcal{A}})$ and that σ converges to zero.*

Justification. The short vectors $\mathbf{v}$ considered will be obtained from random variables $v_{1,1}, \ldots, v_{\rho_K, d_K}$ independently sampled from a centred distribution of

variance σ on $\mathbb{Z}$ such that

$$\|\mathbf{v}\|^2 = \sum_{k=1}^{d_K\rho_K/\delta_{\mathcal{A}}} \left(\sum_{i=1}^{\delta_{\mathcal{A}}} v_{k,i}^2 + \sum_{i\neq j\in[1;\delta_{\mathcal{A}}]} \frac{z_{i,j}}{\delta_{\mathcal{A}}} \cdot v_{k,i}v_{k,j} \right)$$

for some constant $(z_{i,j})_{i,j}$ in $\mathbb{C}^{d_k\times d_K}$ of norm one, and σ converging to a finite limit when $d_K\rho_K$ grows. Then, using the central limit theorem, the delta method provides $\mathbb{E}\left[\ln\|\mathbf{v}\|\right] = \frac{1}{2}\ln(\rho_K d_K) + \ln(\sigma) + O\left(\frac{1}{d_K\rho_K}\right)$.

Second, we model L as constant because we will consider if d_K grows families of fields K with such a common subfield. We model the discriminant term $\ln(D_{\mathcal{A}}^{1/d_K}) - \ln(\delta_{\mathcal{A}})$ as constant even if d_K grows because this term describes a relative density of the field K's (or L's if $\mathcal{A}$ is run on $\kappa_L(M)$) ring of integers with respect to its degree. For cyclotomic number fields, this term only depends on the odd primes in the radical of the conductor, staying constant even if their valuation in the conductor grows.

Third, the success condition described by relation 1 is made under the module-GSA heuristic, which is believed to be a valid approximation if $\rho_K d_K > 2\beta_{\mathcal{A}}$. Denoting $C_L = \frac{1}{2d_L}\ln\left(\frac{|\Delta_L|}{(2\pi e d_L)^{d_L}}\right)$, a module-GSA and module-Gaussian heuristic provide [DEP25]

$$\frac{1}{2}\cdot\mathbb{Q}\text{-slope}_L(\beta_{\mathcal{A}}) = -\frac{1}{\beta_{\mathcal{A}}-d_L}\left(\frac{\ln\beta_{\mathcal{A}}}{2} + C_L + o(1)\right),$$

hence, denoting $C_{\mathcal{A}} = \frac{1}{2\delta_{\mathcal{A}}}\ln\left(\frac{D_{\mathcal{A}}}{(2\pi e\delta_L)^{\delta_L}}\right)$, equality in relation 1 provides

$$\begin{aligned} d_K\rho_K &= \beta_{\mathcal{A}} + \frac{\beta_{\mathcal{A}}-d_L}{\frac{\ln\beta_{\mathcal{A}}}{2} + C_L + o(1)}\left(\frac{\ln(Q/\sigma^2)}{2} + C_{\mathcal{A}} + o(1)\right) \\ \frac{d_K\rho_K}{\beta_{\mathcal{A}}} &= \frac{\ln(Q/\sigma^2)}{\ln\beta} + 1 + o\left(\frac{\ln(Q/\sigma^2)}{\ln\beta}\right) \end{aligned}$$

As $\rho_K d_K$ should be greater than $2\beta_{\mathcal{A}}$, Q/σ^2 should then be an $\Omega(\beta_{\mathcal{A}})$.

Finally, the (module-)GSA heuristic under which relation 1 holds requires not lying in an overstretched regime in which Q would be an $\Omega(\beta_{\mathcal{A}})$, so we require $Q = o(\beta_{\mathcal{A}})$ in the asymptotic analysis of the parameters yielding equality in 1. As $Q = \Omega(\sigma^2\beta_{\mathcal{A}})$, this implies that σ converges to zero. □

5.2 Asymptotic Comparisons of Primal Attacks

Lemma 4 (Heuristic Asymptotic Comparison of Attacks). *For $\mathcal{A}$ and $\mathcal{B}$ two attacks against common uSVP instances, corresponding to algorithm lines in Table 1, with minimal blocksizes over $\mathbb{Q}$, $\beta_{\mathcal{A}}$ and $\beta_{\mathcal{B}}$, provided by an equality in relation 1, under the asymptotic modelling of Heuristic 1,*

$$\beta_{\mathcal{A}} = \beta_{\mathcal{B}} + \frac{\beta_{\mathcal{B}}}{1+\frac{\ln\beta_{\mathcal{B}}}{\ln U}}\left(\frac{\ln\left(\frac{|\Delta_{L_{\mathcal{A}}}|^{1/d_{L_{\mathcal{A}}}d_{L_{\mathcal{B}}}}}{|\Delta_{L_{\mathcal{B}}}|^{1/d_{L_{\mathcal{B}}}d_{L_{\mathcal{A}}}}}\right)}{\ln\beta_{\mathcal{B}}} + \frac{\ln\left(\frac{D_{\mathcal{B}}^{1/d_K}\delta_{\mathcal{A}}}{D_{\mathcal{A}}^{1/d_K}\delta_{\mathcal{B}}}\right)}{\ln U} + o\left(\frac{1}{\ln\beta_{\mathcal{B}}}\right)\right),$$

where $D_{\mathcal{A}}, \delta_{\mathcal{A}}, D_{\mathcal{B}}, \delta_{\mathcal{B}}$ correspond to attacking parameters in Table 1, $\mathcal{A}$ runs the attack on a subfield $L_{\mathcal{A}}$ of K of degree $d_{L_{\mathcal{A}}}$, and respectively for $\mathcal{B}$, and U denotes Q/σ^2, where the $d_K\rho_K$-th root of the volume of the instances over $\mathbb{Q}$ are given by $\sqrt{D_{\mathcal{A}}^{1/d_K}Q}$ for $\mathcal{A}$, $\sqrt{D_{\mathcal{B}}^{1/d_K}Q}$ for $\mathcal{B}$, and σ is the parameter in the distribution of the short vector presented in Heuristic 1. Note that this modelling provides the existence of a bound b such that for every $\beta_{\mathcal{B}} > b$, $U > \beta_{\mathcal{B}}$.

Proof. Let $\mathcal{A}$ and $\mathcal{B}$ be two attack strategies against common uSVP instances, corresponding to respective expected minimal blocksizes over $\mathbb{Q}$ $\beta_{\mathcal{A}}$ and $\beta_{\mathcal{B}}$, provided by equality in relation 1.

When $\beta_{\mathcal{B}}$ grows, Heuristic 1 yields, denoting for a field L, $C_L = \frac{1}{2d_L}\ln\left(\frac{|\Delta_L|}{(2\pi e d_L)^{d_L}}\right)$, and for an attack algorithm $\mathcal{C}$ using L, $C_{\mathcal{C}} = \frac{1}{2d_K}\ln\left(\frac{D_{\mathcal{C}}}{(2\pi e\delta_L)^{d_K}}\right)$,

$$\begin{aligned} d_K\rho_K &= \beta_{\mathcal{B}} + \frac{\beta_{\mathcal{B}} - d_L}{\frac{\ln\beta_{\mathcal{B}}}{2} + C_L + o(1)}\left(\frac{\ln(Q/\sigma^2)}{2} + C_{\mathcal{B}} + o(1)\right) \\ &= \frac{\beta_{\mathcal{B}}\ln(Q/\sigma^2)}{\ln\beta_{\mathcal{B}}} + \beta_{\mathcal{B}} - 2C_L\frac{\beta_{\mathcal{B}}\ln(Q/\sigma^2)}{(\ln\beta_{\mathcal{B}})^2} + 2C_{\mathcal{B}}\frac{\beta_{\mathcal{B}}}{\ln\beta_{\mathcal{B}}} + o\left(\frac{\beta_{\mathcal{B}}\ln(Q/\sigma^2)}{(\ln\beta_{\mathcal{B}})^2}\right), \end{aligned}$$

as well as the same relation on variables corresponding to $\mathcal{A}$, from which one may write $\beta_{\mathcal{A}} = \beta_{\mathcal{B}} + A = \beta_{\mathcal{B}} + o(\beta_{\mathcal{B}})$, and, denoting $U = Q/\sigma^2$,

$$\begin{aligned} A\left(\frac{\ln U}{\ln\beta_{\mathcal{B}}} + 1\right) - 2C_{L_{\mathcal{A}}}\frac{\beta_{\mathcal{B}}\ln U}{(\ln\beta_{\mathcal{B}})^2} + 2C_{\mathcal{A}}\frac{\beta_{\mathcal{B}}}{\ln\beta_{\mathcal{B}}} \\ = -2C_{L_{\mathcal{B}}}\frac{\beta_{\mathcal{B}}\ln U}{(\ln\beta_{\mathcal{B}})^2} + 2C_{\mathcal{B}}\frac{\beta_{\mathcal{B}}}{\ln\beta_{\mathcal{B}}} + o\left(\frac{\beta_{\mathcal{B}}\ln U}{(\ln\beta_{\mathcal{B}})^2}\right), \\ A = \frac{\beta_{\mathcal{B}}}{1 + \frac{\ln\beta_{\mathcal{B}}}{\ln U}}\left(\frac{2(C_{L_{\mathcal{A}}} - C_{L_{\mathcal{B}}})}{\ln\beta_{\mathcal{B}}} + \frac{2(C_{\mathcal{B}} - C_{\mathcal{A}})}{\ln U} + o\left(\frac{1}{\ln\beta_{\mathcal{B}}}\right)\right), \end{aligned}$$

in which there exists a bound b such that for every $\beta_{\mathcal{B}} > b$, $U > \beta_{\mathcal{B}}$ (from Heuristic 1). □

The best attack $\mathcal{A}$ will, from the expression of A, minimise $\frac{C_{L_{\mathcal{A}}}}{\ln\beta_{\mathcal{B}}} - \frac{C_{\mathcal{A}}}{\ln U}$. If $\mathcal{A}$ uses mBKZ, we will show in the following subsection that the optimal parameters in Table 1 correspond to running it on $\kappa_L(M)$, for L a density-preserving (see Definition 8) subfield of K, maximising (the negative) $\mathbb{Q}$-slope$_L(\beta_{\mathcal{A}})$. If $\mathcal{B}$ uses BKZ, then $C_{\mathbb{Q}} = 0$, and the best strategy to maximise $\frac{C_{\mathcal{B}}}{\ln U}$ will be to run the attack on $\kappa_{\mathbb{Q}}(M)$, or, if the conductor of K is a power of two, $\kappa_L(\mathbb{Q})$ for any cyclotomic subfield L of K with a power-of-two conductor.

Definition 7 (Relative Discriminant Term). *For number fields M and N, of respective degrees d_M, d_N, and discriminants Δ_M, Δ_N $e^{2(C_M - C_N)} = \frac{d_N}{d_M}\frac{|\Delta_M|^{1/d_M}}{|\Delta_N|^{1/d_N}}$, and in particular if N is a subfield of M, then*

$$e^{2(C_M - C_N)} = \frac{\mathrm{N}_{N/\mathbb{Q}}(\Delta_{M/N})^{\frac{1}{d_M}}}{[M:N]} = \frac{|\Delta_M|^{1/d_M}}{[M:N]|\Delta_N|^{/d_N}}.$$

We will denote $C_M - C_N$ as $C_{M/N}$, and call it the relative discriminant term.

The **discriminant term** of a number field M is $C_{M/\mathbb{Q}} = C_M$, corresponding to the discriminant gap of [DEP25], with properties studied in their Section 4.3. If M is a cyclotomic number field, then C_M is negative or null, depending only on the odd radical of the conductor of M, as the addition of such terms for cyclotomic number fields with conductors the odd primes in the radical (*cf.* [DEP25, Lemma 10]).

Lemma 5 (Discriminant Term of a Cyclotomic Number Field ([DEP25], Lem. 10)). *If K is a cyclotomic field of conductor c and degree d, then its discriminant term C_K is*

$$C_K = \frac{1}{2}\ln\left(\frac{|\Delta_K|^{1/d}}{d}\right) = \frac{1}{2}\sum_{p\in\mathcal{P}_c}\frac{p-2}{p-1}\ln(p) - \ln(p-1) \leq 0,$$

and is null if and only if $\mathrm{rad}(c) = 2$.

Definition 8 (Density-Preserving Subfield). *For K a number field of degree d_K and L one of its subfields, L will be called* density-preserving *iff* $\mathrm{N}_{L/\mathbb{Q}}\left(\Delta_{K/L}\right) = [K:L]^{d_K}$, i.e. $C_{K/L} = 0$.

From Lemma 3, if K and L are cyclotomic, L being a density-preserving subfield of K is equivalent to it being distortion-free, *i.e.* fulfilling condition (i) of Lemma Lemma 3.

5.3 Optimal Attack Strategies, using BKZ or mBKZ

Lemma 6 (Lower Bound on the Relative Discriminant). *For K a degree-d_K number field and L a subfield of K of degree d_L, the relative discriminant term $C_{K/L}$ is greater than or equal to $-\ln([K:L])/2$,* i.e., *$\mathrm{N}_{L/\mathbb{Q}}(\Delta_{K/L}) \geq 1$, which is equivalent to $|\Delta_K|^{1/d_K} \geq |\Delta_L|^{1/d_L}$.*

Proof. This holds because the relative discriminant $\Delta_{K/L}$ is an ideal of $\mathcal{O}_L$.

Theorem 2 (Upper Bound on the Relative Discriminant inside a Cyclotomic Field). *If K is a cyclotomic number field of degree d_K, then for each of its subfields L, $\mathrm{N}_{L/\mathbb{Q}}(\Delta_{K/L}) \leq [K:L]^{d_K}$,* i.e., *$C_{K/L} \leq 0$.*

Proof. See Appendix B.1.

Lemma 7 (Optimal Asymptotic mBKZ Attacks at the First Order). *Asymptotically, under the heuristic modelling 1, the best mBKZ attacks of Table 1 given an input module M over the cyclotomic field $K \simeq \mathbb{Q}(\zeta_c)$ will run on a $\sigma_L(M)$ or $\kappa_L(M)$ for L a density-preserving subfield of K. In particular, a cyclotomic subfield $L \simeq \mathbb{Q}(\zeta_\ell)$ such that $\mathrm{rad}(c)/\mathrm{rad}(\ell) \in \{1;2\}$ fulfils this condition, and other cyclotomic subfields do not.*

Proof. From Lemma 4, if an attack $\mathcal{A}$ uses mBKZ on the subfield L of K, then it will be asymptotically optimal if it minimises

$$\varepsilon = \frac{\ln(|\Delta_L|)/d_L - \ln(d_L)}{\ln \beta_{\mathcal{B}}} - \frac{\ln(D_{\mathcal{A}})/d_K - \ln(\delta_{\mathcal{A}})}{\ln U},$$

where there is a bound $b > 1$ such that for every $\beta_{\mathcal{B}} > b$, $U > \beta_{\mathcal{B}}$, and $(D_{\mathcal{A}}, \delta_{\mathcal{A}})$ is equal to $(|\Delta_K|, d_K)$ if mBKZ runs on $\sigma_L(M)$ and greater than or equal to $(|\Delta_L|^{[K:L]}, d_L)$ if it runs on $\kappa_L(M)$, with equality iff L is distortion-free or of degree one. As, following Theorem 2, for every subfield L, $\frac{|\Delta_K|^{1/d_K}}{d_K} \leq \frac{|\Delta_L|^{1/d_L}}{d_L}$, and running mBKZ on $\kappa_L(M)$ will be at least as good in minimising this quantity as running it on $\sigma_L(M)$, and one should then asymptotically minimise

$$(\ln(|\Delta_L|)/d_L - \ln(d_L)) \left(\frac{1}{\ln \beta_{\mathcal{B}}} - \frac{1}{\ln U} \right),$$

where for $\beta_{\mathcal{B}} > b$, $\frac{1}{\ln \beta_{\mathcal{B}}} - \frac{1}{\ln U} > 0$. As, according to Theorem 2, $\ln(|\Delta_L|)/d_L - \ln(d_L) \geq \ln(|\Delta_K|)/d_K - \ln(d_K)$, this will be achieved if and only if L is density-preserving, and in a regime where the use of $\sigma_L(M)$ or $\kappa_L(M)$ is indifferent in minimising ε. If L is cyclotomic of conductor ℓ, then Lemma 3 shows that being density-preserving is equivalent to having $\mathrm{rad}(c)/\mathrm{rad}(\ell)$ in $\{1; 2\}$, denoting c the conductor of K, or saying that L is distortion-free (see Theorem 1). □

Lemma 8 (Optimal Asymptotic mBKZ Attacks at the Second Order). *Asymptotically, under the heuristic modelling 1, the best mBKZ attacks of Table 1 given an input module M over the cyclotomic field $K \simeq \mathbb{Q}(\zeta_c)$ will run on $\sigma_L(M)$ or $\kappa_L(M)$ for L a density-preserving subfield of minimal degree, resulting on a constant blocksize gain with respect to the use of density-preserving subfields of higher degrees.*

Proof. Let $\mathcal{A}$ and $\mathcal{B}$ be attacks such that $C_{L_{\mathcal{A}}} = C_{\mathcal{A}} = C_{L_{\mathcal{B}}} = C_{\mathcal{B}} = C$ (obtained from Lemma 7 with the attack strategies for which the statement must be proven, for which $C = C_K$). Equality in relation 1 provides

$$d_K \rho_K = \beta_{\mathcal{A}} + (\beta_{\mathcal{A}} - d_L) \cdot \frac{\frac{\ln(Q/\sigma^2)}{2} + C + \frac{\ln \beta_{\mathcal{A}}}{2\beta_{\mathcal{A}}} + o\left(\frac{\ln \beta_{\mathcal{A}}}{2\beta_{\mathcal{A}}}\right)}{\frac{\ln \beta_{\mathcal{A}}}{2} + C + \frac{\ln \beta_{\mathcal{A}}}{2\beta_{\mathcal{A}}} + o\left(\frac{\ln \beta_{\mathcal{A}}}{2\beta_{\mathcal{A}}}\right)},$$

and the same relation in the $\mathcal{B}$ parameters. Then, denoting $\beta_{\mathcal{A}} = \beta_{\mathcal{B}} + A = \beta_{\mathcal{B}} + o(\beta_{\mathcal{B}})$, as $\rho_K d_K$ is the same in both equations, as $\frac{1}{\ln(\beta_{\mathcal{B}}+A)} = \frac{1}{\ln \beta_{\mathcal{B}}}\Big(1 + \frac{A}{\beta_{\mathcal{B}} \ln \beta_{\mathcal{B}}} + o\left(\frac{A}{\beta_{\mathcal{B}} \ln \beta_{\mathcal{B}}}\right)\Big)$, denoting $U = Q/\sigma^2$,

$$\begin{aligned} A\left(1 + \frac{\ln U}{\ln \beta_{\mathcal{B}}} - \frac{2C \ln U}{(\ln \beta_{\mathcal{B}})^2} + \frac{2C}{\ln \beta_{\mathcal{B}}} - \frac{4C}{(\ln \beta_{\mathcal{B}})^2}\right) \\ = (d_L - d_K)\frac{\ln U}{\ln \beta_{\mathcal{B}}} + o\left(\frac{\ln U}{\ln \beta_{\mathcal{B}}}\right) + o(1) \\ A = \frac{d_{L_{\mathcal{A}}} - d_{L_{\mathcal{B}}} + o(1)}{1 + \frac{\ln \beta_{\mathcal{B}}}{\ln U}}, \end{aligned}$$

and

$$\beta_{\mathcal{A}} = \beta_{\mathcal{B}} + \frac{d_{L_{\mathcal{A}}} - d_{L_{\mathcal{B}}} + o(1)}{1 + \frac{\ln \beta_{\mathcal{B}}}{\ln U}},$$

hence the $L_{\mathcal{A}}$ minimising the blocksize will have the smallest degree, with a constant blocksize difference with the use of K, when running mBKZ on M directly. □

Lemma 9 (Optimal Asymptotic BKZ Attacks). *Asymptotically, under the heuristic modelling 1, the best BKZ attacks of Table 1 given an input module M over the cyclotomic field K will run on $\kappa_L(M)$ minimising $\ln(|\Delta_L|)/d_L - \ln(d_L)$. Considering only cyclotomic subfields L of K, then these will be the ones with conductor a power of two: in particular, running BKZ on $\kappa_{\mathbb{Q}}(M)$ fulfils this condition, with the same blocksize as another BKZ algorithm on a $\kappa_L(M)$ for L a cyclotomic different from $\mathbb{Q}$, up to an $o(1)$ difference.*

Proof. From Lemma 4, if an attack $\mathcal{B}$ uses BKZ, then as $\Delta_{\mathbb{Q}} = d_{\mathbb{Q}}^{d_{\mathbb{Q}}} = 1$, the best asymptotic strategy will maximise $\frac{D_{\mathcal{B}}^{1/d_K}}{\delta_{\mathcal{B}}}$. From Theorem 2, for any subfield L of K, $\frac{|\Delta_K|^{1/d_K}}{d_K} \leq \frac{|\Delta_L|^{1/d_L}}{d_L} \leq \frac{D_{\mathcal{B}}^{1/d_K}}{\delta_{\mathcal{B}}}$, and for cyclotomic L, these two inequalities become equalities iff L is distortion-free (see Lemmas 2, 3 and Theorem 1). If L is cyclotomic, the best strategy for $\mathcal{B}$ will be to run $\mathcal{B}$ on κ_L for $L \simeq \mathbb{Q}(\zeta_\ell)$ a cyclotomic subfield maximising the discriminant term $\frac{|\Delta_L|^{1/d_L}}{d_L} = \prod_{p|\ell} \frac{p^{1-\frac{1}{p-1}}}{p-1} \leq 1$. [DEP25, Section 4.3] shows that the $|\Delta_L| = d_L^{d_L}$ equality is attained if and only if L's conductor is a power of two. In particular, $L = \mathbb{Q}$ is always an optimal choice. The proof of 8 then shows that all such BKZ algorithms result in the same blocksize up to an $o(1)$ difference. □

Remark 3. All the relative embeddings $\kappa_L(M)$ for which L is a power-of-two cyclotomic field have the same canonical embedding into $\mathbb{Q}$ up to a scaling factor, hence corresponding BKZ attacks will actually have the same performance.

Theorem 3 (Asymptotic Gains of mBKZ with Respect to BKZ). *Asymptotically, under the heuristic modelling 1, the best BKZ and mBKZ attacks of Table 1 given an input module M over the cyclotomic field K, corresponding to respective minimal blocksizes over $\mathbb{Q}$ β_{BKZ} and β_{mBKZ}, verify*

$$\beta_{\mathrm{mBKZ}} = \beta_{\mathrm{BKZ}} + \ln\left(\frac{|\Delta_K|^{1/d_K}}{d_K}\right) \cdot \frac{\beta_{\mathrm{BKZ}}}{1 + \frac{\ln \beta_{\mathrm{BKZ}}}{\ln U}} \left(\frac{1}{\ln \beta_{\mathrm{BKZ}}} - \frac{1}{\ln U}\right) + o\left(\frac{\beta_{\mathrm{BKZ}}}{\ln \beta_{\mathrm{BKZ}}}\right)$$

where U denotes Q/σ^2, where $\det_{\mathbb{Q}}(M) = \sqrt{|\Delta_K|^{\rho_K} Q^{\rho_K d_K}}$, and σ is the parameter in the distribution of the short vector of a uSVP instance presented in Heuristic 1, which in particular supposes that there is a bound b such that for every $\beta_{\mathrm{BKZ}} > b$, $U > \beta_{\mathrm{BKZ}}$.

Proof. This directly follows from Lemmas 9 and 7.

6 Application to NTRU+

The South Korean standard for key encapsulation NTRU+KEM [KP25], the related public-key encryption scheme NTRU+PKE [KP24], and some parameter sets for the signature scheme NTRU+Sign [HWK+25], rely on the NTRU problem [HPS98] instantiated over cyclotomic fields $K \simeq \mathbb{Q}(\omega_c)$ with conductor divisible by a power of 3. The corresponding NTRU lattices are hence a natural target for structured lattice reduction over the subfield $L \simeq \mathbb{Q}(\omega_3)$. However, the secret vectors are short in the coefficient geometry on K, rather than the canonical geometry, which somewhat counteracts the improved slope of module lattice reduction. In this section we concretely estimate the blocksize required by Module-BKZ, and compare with that of BKZ applied to the coefficient embedding; we conclude there is only a marginal improvement, reducing the security of NTRU+ keys by at most a few bits in each parameterisation.

6.1 NTRU+

Let $K = \mathbb{Q}[X]/(X^n - X^{n/2} + 1)$ be the NTRU+ field, parameterised by $n = 2^i 3^j$ for some positive integers i, j. K is a cyclotomic field of degree n and conductor

$$c = \begin{cases} 2^i 3^{j+1} & \text{if } i \geq 2 \\ 3^{j+1} & \text{if } i = 1. \end{cases}$$

We let q be a positive integer and define $\mathcal{O}_{K,q} = \mathcal{O}_K/q\mathcal{O}_K$. We write $\mathcal{D}_n$ for the distribution of $a_0 + a_1 X + \ldots + a_{n-1}X^{n-1} \pmod q$ where $a_i \leftarrow \psi_1$ are independent centred binomials.

The security of NTRU+ is estimated by the complexity of solving the following Search-NTRU problem, which corresponds to secret key recovery.[3]

Definition 9 (Search-$\mathrm{NTRU}^+_{n,q}$). *For $f \sim \mathcal{D}_n$, $g \sim \mathcal{D}_n$, and $h = gf^{-1} \in \mathcal{O}_{K,q}$. The Search-$\mathrm{NTRU}^+_{n,q}$ problem is the task of recovering f and g, given an instance (h, n, q).*

For an NTRU+ instance (h, n, q) we define the corresponding $\mathcal{O}_K$-module by its publicly known pseudobasis

$$M_K = \begin{pmatrix} q \\ 0 \end{pmatrix} \mathcal{O}_K + \begin{pmatrix} h \\ 1 \end{pmatrix} \mathcal{O}_K$$

which contains the secret submodule $(g, f)^t \mathcal{O}_K$.

[3] In reality, f is instead sampled via $f' \leftarrow \mathcal{D}_n$, $f = 3f' + 1$ in the NTRU+KEM specification [KP25], and $f' \leftarrow \mathcal{D}_n$, via $f = f' + 2^{-1}$ in NTRU+Sign. However in both cases the security estimates therein are based on the simplified distribution we state here, and in any case such variations are closely related, see [ACD+18, §2.2].

6.2 Concrete Security Estimation

The Two Attacks. Let us analyse unstructured and structured attacks, that Section 5 deemed asymptotically optimal.

Unstructured attack. The standard approach to the search problem is to view M_K as an integer lattice by the coefficient embedding, and then apply BKZ; in fact, our analysis in Section 5 shows that this approach is asymptotically optimal (see Lemma 9, under the modelling of Heuristic 1). From the $\mathcal{O}_K$-basis defining M_K, one attains the integer basis

$$\mathbf{B} = \begin{pmatrix} \mathrm{rot}(q) & \mathrm{rot}(h) \\ \mathrm{rot}(0) & \mathrm{rot}(1) \end{pmatrix} = \begin{pmatrix} q\mathbf{I}_n & \mathrm{rot}(h) \\ \mathbf{0} & \mathbf{I}_n \end{pmatrix} \in \mathbb{Z}^{2n \times 2n}.$$

and the lattice $\Lambda = \Lambda(\mathbf{B})$. Writing $g = \sum g_i X^i$ and $f = \sum f_i X^i$, the lattice contains the short vector $\mathbf{v} = (g_0, \ldots, g_{n-1}, f_0, \ldots, f_{n-1})$ with entries sampled from ψ_1, and thus $\|\mathbf{v}\|^2 \sim \mathrm{Bin}(2n, 1/2)$ with $\mathbb{E}[\|\mathbf{v}\|^2] = n$.

Structured attack. Following analysis in previous sections, we consider the optimal subfield $L = \mathbb{Q}[Y]/(Y^2+Y+1) \simeq \mathbb{Q}(\omega_3)$ of K, where $Y = -X^{n/2}$. By the relative coefficient embedding $\kappa_L : K \to L^{n/2}$ we obtain an $\mathcal{O}_L$-module $M_L = \kappa_L(M_K)$.

M_L contains a dense submodule $\mathbf{v}\mathcal{O}_L$ where $\mathbf{v} = (\kappa_L(g), \kappa_L(f))$. Each entry of $\mathbf{v} \in L^n$ has the form $a + bY$ where $a, b \leftarrow \psi_1$ independently. In the canonical geometry on L,

$$\|a + bY\|^2_{\mathrm{Tr}_L} = 2(a^2 - ab + b^2) = \begin{cases} 0 & \text{with probability } 1/4, \\ 2 & \text{with probability } 5/8, \\ 6 & \text{with probability } 1/8, \end{cases}$$

which has expectation 2, and thus as expected $\mathbb{E}[\|\mathbf{v}\|^2_{\mathrm{Tr}_L}] = 2n$.

Estimation Methodology. Let us describe several estimation methodologies, offering different precision-cost tradeoffs.

GSA/mGSA Estimates. One may apply BKZ on Λ to recover $\mathbf{v}$. However, as is standard for q-ary lattices, it can be beneficial to instead apply the attack to some projected sublattice $\Lambda_{[n-m:2n]}$ for some $1 \leq m \leq n$ to recover a projection of $\mathbf{v}$, see for example [DW21, §2.4]. The standard methodology [ADPS16] thus estimates the required BKZ blocksize β_{GSA} as the smallest $\beta \geq 45$ satisfying

$$\sqrt{\beta/2} \leq \mathrm{gh}(\beta)^{\frac{2\beta-n-m-1}{\beta-1}} \cdot q^{\frac{m}{n+m}}$$

where we are free to choose $\max(\beta - n, 1) \leq m \leq n$ minimising the right hand side. We note our predicted blocksizes in Table 2; these correlate exactly to the bit-security estimates derived in [KP25].

Similarly, if one applies Module-BKZ to the projected module $(M_L)_{[m:n]}$ for $1 \leq m \leq n/2$, the mGSA success condition [PT26] estimates the required (rational) blocksize as $\beta_{\mathrm{mGSA}} = 2 \cdot \beta_L$ where β_L is the minimal solution to

$$\frac{1}{2}\ln(2\beta_L) \leq \mathrm{lgh}_L(\beta_L) + \frac{\ln \det_{\mathbb{Q}}((M_L)_{[m:n]})}{n+2m} + \frac{\rho_L - \beta_L}{2d_L}\mathrm{slope}_L(\beta_L).$$

and m may be chosen optimally. We compute β_{mGSA} for each NTRU+ parameterisation in Table 2 and compute the gain relative to the unstructured attack.

Simulation-based estimate (uSVP model). For the unstructured attack, we employ the estimator of [PT26, §3], computing the success probability for each β of BKZ-β with $\tau = 10$ tours, for each β. This gives a full success probability curve, though in Table 2 we present an 'expected' blocksize $\mathbb{E}[\beta]$ by treating the curve as a cumulative density function as in [PV21, §4.2].

For the structured attack, we employ the analogous estimator of [PT26, §5.3], computing the success probability for each β_L of mBKZ_{β_L} with $\tau_L = \tau d_L$ tours to balance the number of SVP oracle calls relative to BKZ (see [PT26, Appendix D]). We present an equivalent 'expected' blocksize $\mathbb{E}[d_L\beta_L]$ in Table 2.

We note that these estimators require not only $\mathbb{E}[\|\mathbf{v}\|^2]$, but in fact require the exact distribution of $\|\mathbf{v}\|^2$, which we derived above in each case. Finally, we note that for each blocksize we choose m as prescribed by the GSA/mGSA estimate.

Simulation-based estimate (refined). The previous estimates assume a uSVP instance, though in fact in each case we have several short vectors produced by the action of μ_K on $\mathbf{v}$. Thus, as a further refinement, we employ the model of [PT26, §6], with $|\mu_K|/2$ equally[4] short vectors for an unstructured case, and $|\mu_K|/|\mu_L|$ equally short vectors in the structured case, as justified in Heuristic 5 of Appendix C.

Conclusions. We find across all models that mBKZ provides a consistent but small gain ranging 4–10 blocksizes across various parameter sets. In the Core-SVP hardness model [ADPS16], a blocksize β corresponds to 0.292β bits of classical security [BDGL16], thus mBKZ reduces the security by no more than 1–3 bits for any parameterisation.

We remark that in almost every case the simulation-based estimates give higher blocksizes than the (m)GSA estimates, a phenomenon also observed and explained in [PV21, §6.1].

We stress that we only considered Search-NTRU corresponding to the *key recovery attack*, whereas for example the security of $\mathsf{NTRU} + \mathsf{Sign}$ also depends on estimating the *signature forgery* attack, which can be reduced to the Module-SIS problem. Thus, estimating the hardness of Module-SIS would be a natural consideration for future work.

[4] We note this is optimistic for the unstructured case, as the action of $\omega \in \mu_K$ is not necessarily an isometry in the coefficient geometry, though generally does not increase a vector's length too drastically.

Scheme	(m)GSA			Simulation (uSVP)			Simulation (refined)		
	BKZ	mBKZ	Gain	BKZ	mBKZ	Gain	BKZ	mBKZ	Gain
NTRU+576-{KEM,PKE}	396	392	4	403.0	395.4	7.6	394.2	388.2	6.0
NTRU+768-{KEM,PKE}	563	558	5	574.7	567.4	7.3	565.0	559.4	5.7
NTRU+864-{KEM,PKE}	649	644	5	662.4	655.3	7.1	652.3	646.9	5.5
NTRU+1152-{KEM,PKE}	910	906	4	930.9	924.7	6.2	919.8	915.2	4.5
NTRU+648-SIGN	418	412	6	425.8	416.6	9.2	417.1	409.6	7.6
NTRU+972-SIGN	663	656	7	677.4	667.2	10.2	667.5	659.0	8.5
NTRU+1296-SIGN	933	926	7	955.7	945.5	10.3	944.7	936.3	8.5

Table 2: Estimated BKZ blocksizes, compared to mBKZ equivalent blocksizes, required for NTRU+ key recovery. See §6.2 for a detailed explanation. The reported gain is simply the difference between the BKZ and mBKZ blocksizes.

7 Overstretching

Let $K = \mathbb{Q}[X]/\Phi_c(X) \simeq \mathbb{Q}(\omega_c)$ and, following standard cryptographic notation, let $n = d_K$. We consider the NTRU module

$$M_K = \begin{pmatrix} q \\ 0 \end{pmatrix} \mathcal{O}_K + \begin{pmatrix} h \\ 1 \end{pmatrix} \mathcal{O}_K$$

containing the dense submodule $S_K = \binom{g}{f}\mathcal{O}_K = \mathbf{v}\mathcal{O}_K$. Their corresponding volumes are

$$\begin{aligned} \det_{\mathbb{Q}}(M_K) &= |\Delta_K| q^n \\ \det_{\mathbb{Q}}(S_K) &= |\Delta_K|^{1/2} \mathrm{N}(\mathbf{v}) \end{aligned}$$

where as usual $\mathrm{N}(\mathbf{v}) = \mathrm{N}(f\bar{f} + g\bar{g})^{1/2}$.

We let $L \subset K$ be an arbitrary cyclotomic subfield and define $s = [K : L]$. Then the canonical descent[5] $M = \sigma_L(M_K)$ is an $\mathcal{O}_L$-module of rank $2s$, with dense submodule $S = \sigma_L(S_K)$ of rank s. In particular S is generated by the s L-linearly independent short vectors $\{\omega_c^k \mathbf{v}\}_{k=0}^{s-1}$, which all have equal canonical length. We consider reducing M with mBKZ_L, until either of the two following events happens:

- SKR (Secret Key Recovery): a vector $\mathbf{w} \in S$ with the same length as $\mathbf{v}$ is inserted into the pseudobasis;
- DSD (Dense Submodule Discovery): a vector $\mathbf{w} \in S$, strictly longer than $\mathbf{v}$, is inserted into the pseudobasis.

The *fatigue point* is the smallest modulus q for which, when reducing with mBKZ, DSD occurs first more often; for larger q, the mGSA heuristic for SKR becomes a significant overestimate for the security of keys.

7.1 Module Pataki–Tural Lemma

The overstretching analysis of Kirchner and Fouque [KF17] is based on the Pataki–Tural Lemma [PT08] which constrains the shape of a basis profile of a lattice containing dense sublattices. An algebraic analogue for $\mathcal{O}_L$-modules was shown in [KK24], and admits the following clean restatement in terms of the L-profile.

Theorem 4 ([KK24, Thm. 1]). *Let $(\mathbf{b}_i, \mathfrak{J}_i)_{i=1}^{\rho}$ be a pseudobasis of an $\mathcal{O}_L$-module M, with L-profile $(\ell_i)_{i=1}^{\rho}$. For every submodule $S \subset M$ of rank s,*

$$\ln \det_{\mathbb{Q}}(S) \geq \min_{\substack{I \subset \{1,\dots,\rho\} \\ |I|=s}} \sum_{i \in I} \ell_i.$$

[5] Which can be computationally realised by the relative coefficient embedding κ_L (up to a scale factor) when K/L is a distortion-free extension.

Remark 4. Recalling that L-profile may be written in terms of the $\mathbb{Q}$-profile $(\ell_i^{\mathbb{Q}})_{i=1}^{\rho d_L}$ via summing blocks

$$\ell_i = \sum_{j=1}^{d_L} \ell^{\mathbb{Q}}_{(i-1)d_L+j}$$

we may interpret Theorem 4 as a blockwise version of the usual Pataki–Tural lemma, where we only minimise over subsets of the $\mathbb{Q}$-profile which either entirely include or entirely exclude these d_L-sized blocks.

To improve both the asymptotic and concrete analysis of [KF17], Ducas and van Woerden gave a generalisation of the above [DW21, Lem. 3.3]. We state the module analogue below, and prove it using the theory of module filtrations developed in [MS20]. In particular, specialising to $r = 0$ below provides an alternative proof of Theorem 4.

Theorem 5. *Let $(\mathbf{b}_i, \mathfrak{J}_i)_{i=1}^{\rho}$ be a pseudobasis of an $\mathcal{O}_L$-module M, with L-profile $(\ell_i)_{i=1}^{\rho}$. For any rank s submodule $S \subset M$ and index $0 \le r \le \rho$,*

$$\ln \det_{\mathbb{Q}}(S \cap M_{[1:r]}) \le \ln \det_{\mathbb{Q}}(S) - \min_{\substack{I \subset \{r+1,\dots,\rho\} \\ |I| = s-k}} \sum_{i \in I} \ell_i$$

where $k = \operatorname{rank}_L(M_{[1:r]} \cap S)$.

Proof. The pseudobasis of M provides a filtration with tail

$$M_{[1:r]} \subsetneq M_{[1:r+1]} \subsetneq \cdots \subsetneq M_{[1:\rho]} = M$$

where each rank-one projection $\widetilde{M_i} = \pi_i(M_{[1:i]}) = \mathbf{b}_i^* \mathfrak{J}_i$ has volume exactly $\det_{\mathbb{Q}}(\widetilde{M_i}) = e^{\ell_i}$.

For each i define $V_i = \operatorname{span}_{L_{\mathbb{R}}}(M_{[1:i]})$, so that $\pi_i = \pi_{(V_{i-1})^\perp}$. Consider the chain of modules $\Gamma_i = S \cap M_{[1:i]} = S \cap V_i$ obtained by intersecting the filtration tail above with S,

$$(S \cap M_{[1:r]}) = \Gamma_r \subset \Gamma_{r+1} \subset \cdots \subset \Gamma_\rho = S.$$

This chain is in general not yet a filtration tail for S, but we will extract one.

First, for each $i = r+1, \dots, \rho$ the restriction of π_i to Γ_i is an $L_{\mathbb{R}}$-linear map with kernel $\Gamma_i \cap V_{i-1} = \Gamma_{i-1}$, and image $\pi_i(\Gamma_i) \subset \pi_i(M_{[1:i]}) = \widetilde{M_i}$ of rank at most 1. Thus each Γ_{i-1} is a primitive submodule of S, and

$$\operatorname{rank}_L(\Gamma_i) \le \operatorname{rank}_L(\Gamma_{i-1}) + 1$$

or in other words, the rank increases by at most 1 at each step. It begins at $\operatorname{rank}_L(S \cap M_{[1:r]}) = k$ by definition, and ends at $\operatorname{rank}_L(S) = s$. If the rank does not increase from index $i-1$ to i then in fact $\Gamma_{i-1} = \Gamma_i$, being primitive submodules of S with the same $L_{\mathbb{R}}$-span. Thus the chain is in fact a filtration, up to potential duplicates.

To remedy this, define the map $\iota : \{0, \ldots, s-k\} \to \{r, \ldots, \rho\}$ where $\iota(j)$ is the minimal index with $\mathrm{rank}_L(\Gamma_{\iota(j)}) = k + j$. By the discussion above, this gives a filtration tail for S,

$$S \cap M_{[1:r]} = \Gamma_{\iota(0)} \subsetneq \Gamma_{\iota(1)} \subsetneq \cdots \subsetneq \Gamma_{\iota(s-k)} = S$$

and hence by Lemma 1,

$$\det_{\mathbb{Q}}(S) = \det_{\mathbb{Q}}(S \cap M_{[1:r]}) \prod_{j=1}^{s-k} \det_{\mathbb{Q}}(\widetilde{\Gamma}_{\iota(j)}) \tag{2}$$

for the rank-one projections $\widetilde{\Gamma}_{\iota(j)} = \pi_{(\Gamma_{\iota(j-1)})^\perp}(\Gamma_{\iota(j)})$. By minimality of $\iota(j)$ we have $\mathrm{rank}_L(\Gamma_{\iota(j)-1}) = k + j - 1$, so by the discussion above the restriction of $\pi_{\iota(j)}$ to $\Gamma_{\iota(j)}$ has image of rank exactly 1. Moreover $\Gamma_{\iota(j-1)} \subset V_{\iota(j)-1}$, so that $\pi_{\iota(j)} = \pi_{\iota(j)} \circ \pi_{(\Gamma_{\iota(j-1)})^\perp}$, and hence

$$\pi_{\iota(j)}(\widetilde{\Gamma}_{\iota(j)}) = \pi_{\iota(j)}(\Gamma_{\iota(j)}) \subset \pi_{\iota(j)}(M_{[1:\iota(j)]}) = \widetilde{M}_{\iota(j)}$$

is an inclusion of rank-1 $\mathcal{O}_L$-modules. We see

$$\det_{\mathbb{Q}}(\widetilde{\Gamma}_{\iota(j)}) \geq \det_{\mathbb{Q}}(\pi_{\iota(j)}(\widetilde{\Gamma}_{\iota(j)})) \geq \det_{\mathbb{Q}}(\widetilde{M}_{\iota(j)}) = e^{\ell_{\iota(j)}} \tag{3}$$

where the first inequality is because the $L_{\mathbb{R}}$-linear projection $\pi_{\iota(j)}$ cannot increase the volume [MS20, Cor. 3.4.4], and the second is because any full-rank sublattice has determinant at least that of the full lattice. The result follows by combining Equations (2) and (3), taking logs, and defining the index set $I = \{\iota(1), \ldots, \iota(s-k)\}$.

7.2 Heuristics

To enable both asymptotic and concrete analysis we assume the following heuristics, generalising those in prior work to modules over other fields than the rationals.

Heuristic 2 (Rank Heuristic). *Let $M = \sigma_L(M_K)$ be the canonical descent of an NTRU module, with rank $2s$ and dense submodule S of rank s. Then at every step during lattice reduction before the DSD event is triggered, we assume $\mathrm{span}_{L_{\mathbb{R}}}(M_{[1:s+k]})$ and $\mathrm{span}_{L_{\mathbb{R}}}(S)$ behave as randomly oriented subspaces of rank $s+k$ and s respectively. Hence, we assume $\mathrm{rank}_L(M_{[1:s+k]} \cap S) = k$. Then Theorem 5 specialises to*

$$\ln \det_{\mathbb{Q}}(M_{[1:s+k]} \cap S) \leq \ln \det_{\mathbb{Q}}(S) - \sum_{i=s+k+1}^{2s} \ell_i.$$

Heuristic 3 (DSD-PT estimate). *A tour of $\mathrm{mBKZ}_L^{\beta_L}$ triggers the DSD event if*

$$\|\pi_{s+k-\beta_L+1}(\mathbf{w})\| \leq \lambda_1(M_{[s+k-\beta_L+1\,:\,s+k]})$$

where $\mathbf{w}$ is a shortest vector of $M_{[1:s+k]} \cap S$ for some $1 \leq k \leq s$.

Heuristic 4 (mZGSA). *The public pseudobasis of $M = \sigma_L(M_K)$, after* mBKZ_L *with blocksize β_L, has L-profile $(\ell_i)_{i=1}^{2s}$ of the following shape:*

$$\ell_i = \frac{1}{2s}\ln|\Delta_K| + \begin{cases} d_L \ln q & i \le s + \frac{1}{2} - w \\ \frac{d_L \ln q}{2} + \frac{2s+1-2i}{2}\ln \alpha_L & s + \frac{1}{2} - w < i < s + \frac{1}{2} + w \\ 0 & i \ge s + \frac{1}{2} + w \end{cases}$$

with parameters $w = \frac{d_L \ln q}{2 \ln \alpha_L}$, and $\ln \alpha_L = -\mathrm{slope}_L(\mathrm{mBKZ}_L^{\beta_L})$.

7.3 Asymptotic Analysis

We fix the base field L and consider taking K of large degree $n = d_K = [K : \mathbb{Q}]$, letting $n \to \infty$ in the asymptotic regime

$$\begin{aligned} d_L &= O(1) & \beta &= \mathcal{B}n \\ q &= \Theta(n^{\mathcal{Q}}) & \|\mathbf{v}\| &= O(n^{\mathcal{V}+1/2}) \end{aligned}$$

Remark 5. We stress that $\|\mathbf{v}\|$ above refers to the canonical norm of $\mathbf{v} = (f, g)$. Thus if f and g are sampled as a polynomial in $\mathcal{O}_K$ with short coefficients of variance $\sigma^2 \in \Theta(1)$, then $\mathbb{E}[\|\mathbf{v}\|^2] = 2n^2\sigma^2$ and one sets $\mathcal{V} = 1/2$.

We also note the second condition is equivalent to $\beta_L = \mathcal{B}s$ for $s = [K : L]$, and the final two are equivalent to $q = \Theta(s^{\mathcal{Q}})$ and $\|\mathbf{v}\| = O(s^{\mathcal{V}+1/2})$.

Heuristic Claim 1. *Asymptotically, the mZGSA parameters are*

$$\ln \alpha_L = \frac{d_L^2 \ln n}{\mathcal{B}n} + O\left(\frac{1}{n}\right), \quad w = \frac{\mathcal{Q}\mathcal{B}}{2d_L}n + O\left(\frac{n}{\ln n}\right).$$

Justification. Recall $\ln \alpha_L = -\mathrm{slope}_L(\mathrm{mBKZ}_L^{\beta_L})$ and by the asymptotic analysis in [DEP25, §5] restricted to first-order,

$$\ln(\alpha_L) = \frac{d_L^2}{\beta - d_L}\left(\ln(\beta) + O(1)\right)$$

which when specialised to the regime above gives the result.

We note that these asymptotics are equivalent to the unstructured case; the improved slope of Module-BKZ is invisible at the level of first-order asymptotics. Thus we do not expect to see a gain in this regime, which we confirm with the claim below.

Heuristic Claim 2. *The* $\mathrm{mBKZ}_L^{\beta_L}$ *algorithm with blocksize $\beta_L = \mathcal{B}s$ applied to an NTRU instance over K with parameters $q = \Theta(s^{\mathcal{Q}})$, $\|\mathbf{v}\| = O(s^{\mathcal{V}+1/2})$ triggers the DSD event if*

$$\mathcal{B} = \frac{8\mathcal{V}}{\mathcal{Q}^2 + 1} + o(1).$$

Justification. One follows the proof of [DW21, Claim 3.5], with minor but careful adaptations. The volume of the dense submodule is approximated to first order as $\det_{\mathbb{Q}}(S) = |\Delta_K|^{1/2}\mathrm{N}(\mathbf{v}) \leq |\Delta_K|^{1/2}\|\mathbf{v}\|^n/\sqrt{n^n}$. The structured and unstructured Gaussian heuristics have difference $\mathrm{lgh}_{\mathbb{Q}}(d_L r) - \mathrm{lgh}_L(r) = O(1/r)$, which is asymptotically ignored. The granularity of the module Pataki-Tural lemma is irrelevant as $d_L = O(1)$.

The asymptotic SKR condition similarly hides the gain of Module-BKZ, thus taking $\mathcal{V} = 1/2$ we recover the same asymptotic fatigue point $q = n^{2.484+o(1)}$ as in [DW21], and conclude that Module-BKZ does not provide a gain at the coarse scale of these first order asymptotics.

7.4 Concrete Analysis

The $o(1)$ term in coarse asymptotic analysis above hides many terms, such as the improved slope of mBKZ, which would concretely lower the fatigue point; and in any case the asymptotic formula does not determine whether finite parameter sets lie in the overstretched regime. Thus, we provide a concrete analysis of the fatigue point adapting [DW21, §4] to the module setting, implement this estimator programmatically, and apply it to NTRU+.

Volume of the dense submodule. To proceed we will need a model for the expected log-volume of the dense submodule in the NTRU module M_K. In practice f and g are sampled as $\sum_{i=0}^{n-1} f_i X^i, \sum_{i=0}^{n-1} g_i X^i$, with all f_i, g_i independently drawn from a short distribution $\mathcal{D}$ on $\mathbb{Z}$. Modelling these as continuous Gaussians we obtain the following.

Lemma 10. *If $f_i, g_i \sim \mathcal{N}(0, \sigma^2)$ are all independent, then, denoting $\mathbf{v} = (g_0, \ldots, g_{n-1}, f_0, \ldots, f_{n-1})$,*

$$\mathbb{E}[\log \mathrm{N}(\mathbf{v})] = \frac{n}{2}\left(\ln(n\sigma^2) + \psi(2)\right) - \varepsilon$$

where ε is a correction term depending only on K, made explicit in the proof.

Proof. See Appendix A.4.

Intersection Volumes. For estimation we will also require an average-case model for Theorem 5, quantifying the average slack in the inequality.

Heuristic Claim 3. *Let $M = \sigma_L(M_K)$ be a rank $2s$ NTRU module over $\mathcal{O}_L$, with dense submodule S of rank s and L-profile $(\ell_i)_{i=1}^{2s}$. Under Heuristic 2 we have for $1 \leq k \leq s$,*

$$\mathbb{E}[\ln \det_{\mathbb{Q}}(S \cap M_{[1:s+k]})] = \ln \det_{\mathbb{Q}}(S) - \sum_{j=s+k+1}^{2s} \ell_j + \sum_{j=k+1}^{s} \delta_j$$

where the correction terms are

$$\delta_j = d_L \left[\frac{1}{2}\psi\left(\frac{d_L j}{2}\right) - \frac{1}{2}\psi\left(\frac{d_L(s+j)}{2}\right) + \mathrm{skew}_L(j) - \mathrm{skew}_L(s+j) \right] + \frac{\zeta_L'(j)}{\zeta_L(j)}$$

using the heuristic models for skewness and index gap from [DEP25].

Justification. We follow the proof of Theorem 5, noting that under Heuristic 2 the chain $\Gamma_i = S \cap M_{[1:i]}$ for $i = s+k, \ldots, 2s$ already gives a filtration tail for S with $\mathrm{rank}_L(\Gamma_i) = i - s$ (there is no need for reindexing along a subsequence). The proof is then tight up to the inequalities in Equation (3),

$$\det_{\mathbb{Q}}(\widetilde{\Gamma}_i) \geq \det_{\mathbb{Q}}(\pi_i(\widetilde{\Gamma}_i)) \geq \det_{\mathbb{Q}}(\widetilde{M}_i) = e^{\ell_i}.$$

Write the rank-one projection

$$\widetilde{\Gamma}_i = \mathfrak{a}_i \mathbf{u}_i, \quad \mathbf{u}_i = \pi_{\mathrm{span}_{L_{\mathbb{R}}}(\Gamma_i)}(\mathbf{b}_i^*) \in \mathrm{span}_{L_{\mathbb{R}}}(\Gamma_i) \cap (\Gamma_{i-1})^{\perp},$$

where the inclusion holds since $\mathbf{b}_i^* \perp V_{i-1} \supseteq \Gamma_{i-1}$. Then taking the projection $\pi_i(\widetilde{\Gamma}_i) = \mathfrak{a}_i \pi_i(\mathbf{u}_i)$ and computing determinants,

$$\frac{\det_{\mathbb{Q}}(\widetilde{\Gamma}_i)}{\det_{\mathbb{Q}}(\pi_i(\widetilde{\Gamma}_i))} = \frac{\mathrm{N}(\mathbf{u}_i)}{\mathrm{N}(\pi_i(\mathbf{u}_i))}.$$

Next, $\pi_i(\widetilde{\Gamma}_i) \subset \widetilde{M}_i$ is an inclusion of rank-1 $\mathcal{O}_L$-modules, so there exists an integral ideal $\mathfrak{c} \subset \mathcal{O}_L$ such that $\pi_i(\widetilde{\Gamma}_i)\mathfrak{c}^{-1} = \widetilde{M}_i$ and hence

$$\frac{\det_{\mathbb{Q}}(\pi_i(\widetilde{\Gamma}_i))}{\det_{\mathbb{Q}}(\widetilde{M}_i)} = \mathrm{N}(\mathfrak{c}).$$

Multiplying these two equalities and taking logs,

$$\ln \det_{\mathbb{Q}}(\widetilde{\Gamma}_i) - \ell_i = -\left[\ln \mathrm{N}(\pi_i(\mathbf{u}_i)) - \ln \mathrm{N}(\mathbf{u}_i)\right] + \ln \mathrm{N}(\mathfrak{c}).$$

It remains to model the expectation of these two components; telescoping the sum over $i = s+k+1, \ldots, 2s$ then gives the result. For the first, note that $\langle \mathbf{u}_i, \mathbf{b}_i^* \rangle_L = \langle \mathbf{u}_i, \mathbf{u}_i \rangle_L$, so that $\pi_i(\mathbf{u}_i) = \frac{\langle \mathbf{u}_i, \mathbf{u}_i \rangle_L}{\langle \mathbf{b}_i^*, \mathbf{b}_i^* \rangle_L} \mathbf{b}_i^*$, and therefore

$$\begin{aligned}
&\mathbb{E}[\ln \mathrm{N}(\pi_i(\mathbf{u}_i)) - \ln \mathrm{N}(\mathbf{u}_i)] = \mathbb{E}[\ln \mathrm{N}(\mathbf{u}_i) - \ln \mathrm{N}(\mathbf{b}_i^*)] \\
&= d\left(\mathbb{E}[\ln\|\mathbf{u}_i\| - \ln\|\mathbf{b}_i^*\|] - \mathbb{E}\ln \frac{\sqrt{d}\mathrm{N}(\mathbf{b}_i^*)^{1/d}}{\|\mathbf{b}_i^*\|} + \mathbb{E}\ln \frac{\sqrt{d}\mathrm{N}(\mathbf{u}_i)^{1/d}}{\|\mathbf{u}_i\|} \right) \\
&= d\left(\frac{1}{2}\psi\left(\frac{d(i-s)}{2}\right) - \frac{1}{2}\psi\left(\frac{di}{2}\right) + \mathrm{skew}_L(i-s) - \mathrm{skew}_L(i) \right)
\end{aligned}$$

where the final equality is heuristic: $\mathbf{u}_i$ is the projection of $\mathbf{b}_i^* \in V_i$ onto the rank $i-s$ subspace $\mathrm{span}_{L_{\mathbb{R}}}(\Gamma_i)$, a heuristically random rank $i-s$ subspace of V_i (Heuristic 2), thus we apply [DW21, Lem. 4.1] and the definition of skewness.

For the second, let $(\mathbf{g}_t, \mathcal{I}_t)_{t=1}^{i-s}$ be a pseudobasis of Γ_i and write $\pi_i(\mathbf{g}_t) = x_t \mathbf{b}_i^*$ with $x_t \in L$. Then $\pi_i(\widetilde{\Gamma}_i) = \pi_i(\Gamma_i) = (\sum_t x_t \mathcal{I}_t)\, \mathbf{b}_i^*$, and comparing with $\widetilde{M}_i = \mathfrak{J}_i \mathbf{b}_i^*$ gives

$$\mathfrak{c} = \sum_{t=1}^{i-s} x_t \mathcal{I}_t \mathfrak{J}_i^{-1},$$

the sum of $i-s$ integral ideals. Applying the modelling of the index gap in [DEP25, §4.5], we conclude $\mathbb{E}[\ln \mathrm{N}(\mathfrak{c})] = -\zeta_L'(i-s)/\zeta_L(i-s)$. □

Estimated Fatigue Points. We adapt the fatigue point estimator of [DW21] to the module setting. In short, for each blocksize β_L one assumes the mZGSA (Heuristic 4) and estimates the probability DSD and/or SKR occurs.

Using Heuristic 3, we predict that DSD occurs when, at some index $1 \le k \le s$ we have

$$\|\pi_{s+k-\beta_L+1}(\mathbf{w})\| \le \lambda_1(M_{[s+k-\beta_L+1\,:\,s+k]})$$

where $\mathbf{w}$ is a shortest vector of $M_{[1:s+k]} \cap S$. We estimate the left hand side with the following steps:

- estimate the volume of S using Lemma 10,
- estimate the volume of the intersection $M_{[1:s+k]} \cap S$ using Claim 3
- estimate the length of $\mathbf{w}$ by the Gaussian heuristic gh_L for $M_{[1:s+k]} \cap S$,
- model the length of $\pi(\mathbf{w})$ by an appropriate chi-squared distribution.[6]

We model the right hand side by applying the Gaussian heuristic to the relevant block of the mZGSA profile. We predict that SKR occurs as before, noting that S contains $|\mu_K|/|\mu_L|$ short vectors (up to multiplication by μ_L). We compute the fatigue point q^* as the first value of the modulus q where, when running Progressive-mBKZ, DSD occurs first with probability ≥ 0.5.

Results. We apply the estimator to NTRU+, where $K \simeq \mathbb{Q}(\omega_c)$ for $c = 2^i 3^j$, and f, g are sampled with centred binomial ψ_1 coefficients. We estimate the fatigue point q^*_{mBKZ} for an attacker choosing $L \simeq \mathbb{Q}(\omega_3)$. As a baseline we compare to the fatigue point q^*_{BKZ} for an attacker choosing $L \simeq \mathbb{Q}$, that is to say, applying standard BKZ to the canonical embedding. We give the results for each NTRU+ parameterisation in Table 3, alongside the ratio $q^*_{\mathrm{mBKZ}}/q^*_{\mathrm{BKZ}}$. We conclude that while mBKZ_L reduces the NTRU+ fatigue point by a factor approximately 0.91, none of the parameterisations lie within the overstretched regime.

[6] Replacing this with an analogue of the beta model of [PT26] would likely improve accuracy, but we keep the chi-squared distribution for ease of implementation and consistency with [DW21].

Scheme	q	q^*_{BKZ}	q^*_{mBKZ}	$q^*_{\mathrm{mBKZ}}/q^*_{\mathrm{BKZ}}$
NTRU+576-{KEM,PKE}	3457	19231	17483	0.909
NTRU+768-{KEM,PKE}	3457	39398	35995	0.914
NTRU+864-{KEM,PKE}	3457	52875	48452	0.916
NTRU+1152-{KEM,PKE}	3457	108577	99901	0.920
NTRU+648-Sign	7129	25787	23509	0.912
NTRU+972-Sign	9721	70978	65117	0.917
NTRU+1296-Sign	9721	145825	134527	0.923

Table 3: Estimated fatigue points for the NTRU+ parameterisations, subject to reduction via BKZ ($L \simeq \mathbb{Q}$) and mBKZ ($L \simeq \mathbb{Q}(\omega_3)$), with the ratio between them.

Acknowledgements. Paola de Perthuis was supported by the Dutch Research Council (NWO/OCW) as part of the Quantum Software Consortium programme (project number 024.003.037/3368). All the content of this article is the authors' original work. Artificial intelligence was used to obtain references and background intuition about prior work and the objects studied, and for programming assistance.

References

ABD16. Martin R. Albrecht, Shi Bai, and Léo Ducas. A subfield lattice attack on overstretched NTRU assumptions - cryptanalysis of some FHE and graded encoding schemes. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 153–178. Springer, Berlin, Heidelberg, August 2016.

ACD+18. Martin R. Albrecht, Benjamin R. Curtis, Amit Deo, Alex Davidson, Rachel Player, Eamonn W. Postlethwaite, Fernando Virdia, and Thomas Wunderer. Estimate all the LWE, NTRU schemes! In Dario Catalano and Roberto De Prisco, editors, *SCN 18*, volume 11035 of *LNCS*, pages 351–367. Springer, Cham, September 2018.

ADPS16. Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange - A new hope. In Thorsten Holz and Stefan Savage, editors, *USENIX Security 2016*, pages 327–343. USENIX Association, August 2016.

Apo76. Tom M. Apostol. *Introduction to Analytic Number Theory*. Undergraduate Texts in Mathematics. Springer-Verlag, New York, 1976.

BC22. Iván Blanco-Chacón. On the rlwe/plwe equivalence for cyclotomic number fields. *Applicable Algebra in Engineering, Communication and Computing*, 33(1):53–71, Jan 2022.

BDGL16. Anja Becker, Léo Ducas, Nicolas Gama, and Thijs Laarhoven. New directions in nearest neighbor searching with applications to lattice sieving. In Robert Krauthgamer, editor, *27th SODA*, pages 10–24. ACM-SIAM, January 2016.

BEF+17. Jean-François Biasse, Thomas Espitau, Pierre-Alain Fouque, Alexandre Gélin, and Paul Kirchner. Computing generator in cyclotomic integer rings - A subfield algorithm for the principal ideal problem in $L_{|\Delta_{\mathbb{K}}|}(\frac{1}{2})$ and application to the cryptanalysis of a FHE scheme. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part I*, volume 10210 of *LNCS*, pages 60–88. Springer, Cham, April / May 2017.

BN25. Henry Bambury and Phong Q. Nguyen. Cryptanalysis of an efficient signature based on isotropic quadratic forms. In Ruben Niederhagen and Markku-Juhani O. Saarinen, editors, *Post-Quantum Cryptography - 16th International Workshop, PQCrypto 2025, Part II*, pages 153–175. Springer, Cham, April 2025.

CDPR16. Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 559–585. Springer, Berlin, Heidelberg, May 2016.

CDW17. Ronald Cramer, Léo Ducas, and Benjamin Wesolowski. Short stickelberger class relations and application to ideal-SVP. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part I*, volume 10210 of *LNCS*, pages 324–348. Springer, Cham, April / May 2017.

DD12. Léo Ducas and Alain Durmus. Ring-LWE in polynomial rings. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, volume 7293 of *LNCS*, pages 34–51. Springer, Berlin, Heidelberg, May 2012.

DEP25. Léo Ducas, Lynn Engelberts, and Paola de Perthuis. Predicting module-lattice reduction. In Goichiro Hanaoka and Bo-Yin Yang, editors, *ASIACRYPT 2025, Part III*, volume 16247 of *LNCS*, pages 133–166. Springer, Singapore, December 2025.

DM22. Gabrielle De Micheli and Daniele Micciancio. A fully classical LLL algorithm for modules. Cryptology ePrint Archive, Report 2022/1356, 2022.

DW21. Léo Ducas and Wessel P. J. van Woerden. NTRU fatigue: How stretched is overstretched? In Mehdi Tibouchi and Huaxiong Wang, editors, *ASIACRYPT 2021, Part IV*, volume 13093 of *LNCS*, pages 3–32. Springer, Cham, December 2021.

GM25. Andrea Di Giusto and Chiara Marcolla. Breaking the power-of-two barrier: noise estimation for BGV in NTT-friendly rings. *Designs, Codes and Cryptography*, 93(3):467–502, 2025.

GSVV24. Nihar Gargava, Vlad Serban, Maryna Viazovska, and Ilaria Viglino. Effective module lattices and their shortest vectors. 2024.

HPS98. Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In *Third Algorithmic Number Theory Symposium (ANTS)*, volume 1423 of *LNCS*, pages 267–288. Springer, June 1998.

HWK+25. Ga Hee Hong, Joo Woo, Jonghyun Kim, Minkyu Kim, Hochang Lee, and Jong Hwan Park. More NTRU+sign signatures from cyclotomic trinomials. Cryptology ePrint Archive, Report 2025/612, 2025.

KEF19. Paul Kirchner, Thomas Espitau, and Pierre-Alain Fouque. Algebraic and euclidean lattices: Optimal lattice reduction and beyond. Cryptology ePrint Archive, Report 2019/1436, 2019.

KF17. Paul Kirchner and Pierre-Alain Fouque. Revisiting lattice attacks on overstretched NTRU parameters. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part I*, volume 10210 of *LNCS*, pages 3–26. Springer, Cham, April / May 2017.

KK24. Alexander Karenin and Elena Kirshanova. Finding dense submodules with algebraic lattice reduction. In Serge Vaudenay and Christophe Petit, editors, *AFRICACRYPT 24*, volume 14861 of *LNCS*, pages 403–427. Springer, Cham, July 2024.

KP23. Jonghyun Kim and Jong Hwan Park. NTRU+: Compact construction of NTRU using simple encoding method. *IEEE Transactions on Information Forensics and Security*, 2023. Full version: ePrint 2022/1664.

KP24. Jonghyun Kim and Jong Hwan Park. NTRU+PKE: Efficient public-key encryption schemes from the NTRU problem. Cryptology ePrint Archive, Report 2024/1282, 2024.

KP25. Jonghyun Kim and Jong Hwan Park. NTRU+: Algorithm specification and supporting documentation. Korean Post-Quantum Cryptography Competition (KpqC), 2025.

Kpq25. KpqC Research Group. KpqC algorithms. Korean Post-Quantum Cryptography Research Group (KpqC), 2025. Web page dated 30 January 2026. Accessed 18 October 2026.

LLL82. A. K. Lenstra, Jr. Lenstra, H. W., and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, 1982.

LPSW19. Changmin Lee, Alice Pellet-Mary, Damien Stehlé, and Alexandre Wallet. An LLL algorithm for module lattices. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part II*, volume 11922 of *LNCS*, pages 59–90. Springer, Cham, December 2019.

MS20. Tamalika Mukherjee and Noah Stephens-Davidowitz. Lattice reduction for modules, or how to reduce ModuleSVP to ModuleSVP. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 213–242. Springer, Cham, August 2020.

Neu99. Jürgen Neukirch. *Algebraic Number Theory*, volume 322 of *Grundlehren der mathematischen Wissenschaften*. Springer, Berlin, Heidelberg, 1999.

Pon34. Lev Pontrjagin. The theory of topological commutative groups. *Annals of Mathematics*, 35(2):361–388, 1934.

PT08. Gabor Pataki and Mustafa Tural. On sublattice determinants in reduced bases, 2008.

PT26. Paola de Perthuis and Filip Trenkić. Refined modelling of the primal attack, and variants against module-learning with errors. In *Post-Quantum Cryptology*, April 2026. https://eprint.iacr.org/2025/2195.

PV21. Eamonn W. Postlethwaite and Fernando Virdia. On the success probability of solving unique SVP via BKZ. In Juan Garay, editor, *PKC 2021, Part I*, volume 12710 of *LNCS*, pages 68–98. Springer, Cham, May 2021.

Was97. Lawrence C. Washington. *Introduction to Cyclotomic Fields*, volume 83 of *Graduate Texts in Mathematics*. Springer, New York, 2 edition, 1997.

Appendix

A Distortion between Embeddings

A.1 Distortion when Sampling Spherically in a Cyclic Embedding.

Sampling spherically in the cyclic embedding will not necessarily result in spherical sampling in the Minkowski one, however, we show that it preserves independence of the coefficients of this embedding. The reader may find more extensive discussions on these aspects in [BC22, Section 2.2.3] and [DD12].

For K a cyclotomic number field of conductor c and degree $d = \varphi(c)$, let $A_1, \ldots, A_c$ be independent variables sampled according to a distribution $\mathcal{D}$ over $\mathbb{Z}$, yielding the element $A(X) = \sum_i A_i X^{i-1}$ of $\mathbb{Z}[X]/(X^c - 1)$. For each index j in $[c]$ that is coprime with c, we set, denoting $K \simeq \mathbb{Q}(\zeta_c)$,

$$B_j = \frac{1}{\sqrt{c}} A(\zeta_c^j).$$

The random vector $B = (B_j)_{\substack{j \in [c] \\ j \nmid c}}$ is defined using a Vandermonde matrix:

$$B = \frac{1}{\sqrt{c}} \begin{pmatrix} 1 & \omega_1 & \cdots & \omega_1^{c-1} \\ 1 & \omega_2 & \cdots & \omega_2^{c-1} \\ \vdots & & & \vdots \\ 1 & \omega_d & \cdots & \omega_d^{c-1} \end{pmatrix} \begin{pmatrix} A_1 \\ \vdots \\ A_c \end{pmatrix},$$

where $\omega_1 = \zeta_c, \omega_2, \ldots, \omega_d$ are the Galois conjugates of ω_c.

As the A_i variables are sampled independently, their covariance matrix is the identity times the variance of the distribution $\mathcal{D}$, denoted σ^2. Given j and k in $[c]$ coprime with c,

$$\begin{aligned} B_j \overline{B}_k &= \frac{1}{c} \sum_{i,\ell=1}^{c} A_i \overline{A}_\ell \zeta_c^{j(i-1)} \overline{\zeta}_c^{k(\ell-1)} \\ &= \frac{1}{c} \left(\sum_i A_i^2 (\zeta_c^j \overline{\zeta_c^k})^{i-1} + \sum_{i \neq \ell} A_i A_\ell \zeta_c^{j(i-1)} \zeta_c^{-k(\ell-1)} \right), \end{aligned}$$

and by linearity of the expected value, and as $\zeta_c^j \overline{\zeta_c^k}$ is a c-th root of unity,

$$\mathrm{Cov}(B_a, B_b) = \begin{cases} 0 & \text{if } j \neq k \\ \sigma^2 & \text{otherwise} \end{cases}.$$

Hence, the components of B are also sampled independently, and each has a variance σ^2. Additionally, by linearity, the expected value of B is a null vector.

For B_i one of B's components, the distribution of $B_i\overline{B}_i$ is the one of

$$B_i\overline{B}_i = \frac{1}{c}\left(\sum_{k=1}^{c} A_k^2 + \sum_{k\neq\ell} A_k A_\ell\ \zeta_c^{i(k-\ell)}\right),$$

and the Euclidean norm of B is provided by

$$\|B\| = \frac{1}{\sqrt{c}} \cdot \sqrt{d\left(\sum_{k=1}^{c} A_k^2\right) + \sum_{k\neq\ell} A_k A_\ell \sum_{\substack{j\in[c]\\ j\wedge c=1}} \zeta_c^{j(k-\ell)}}.$$

In particular, note that the expected value of $\|B\|^2$ is $d\sigma^2$.

A.2 Proof of Lemma 2

Proof. As V and N are two $\mathbb{Z}$-bases of the same ring, $\mathcal{O}_L[\zeta_c] = \mathcal{O}_K$, there exists a unimodular matrix $U \in \mathcal{GL}_{d_K}(\mathbb{Z})$ such that $V = NU$, thus M's coefficients are in $\mathbb{Z}$. M is then orthogonal iff it is a signed permutation, *i.e.* iff, denoting $V_1, \ldots, V_{d_K}$ the columns of V and $N_1, \ldots, N_{d_K}$ those of N, there exists a permutation s of $\mathfrak{S}_{d_K}$ and $\varepsilon = (\varepsilon_1, \ldots, \varepsilon_{d_K}) \in \{-1; 1\}^{d_K}$ such that for each index k, $V_k = \varepsilon_k N_{s(k)}$, *i.e.*, as each line corresponds to an embedding of K, $X^{k-1} = \varepsilon_k Y^{i-1} Z^{j-1}$, denoting $s(k) = (j-1)d_L + i$ with $i \in [d_L]$ and $j \in [[K : L]]$. For each such tuple of indices, we will denote $(i, j) = \psi(k)$, where ψ is a bijection from $[d_K]$ to $[d_L] \times [[K : L]]$ corresponding to s.

A. If $d_L \geq 2$.

($\implies$). Let us suppose that there exists ε in $\{-1; 1\}^{d_K}$ and ψ a bijection from $[d_K]$ to $[d_L] \times [[K : L]]$ such that, for each k in $[d_K]$, $X^{k-1} = \varepsilon_k Y^{i-1} Z^{j-1}$, with $(i, j) = \psi(k)$.

In particular, for each i in $[d_L]$, $\varepsilon_{\psi^{-1}(i,1)} X^{\psi^{-1}(i,1)-1}$ is in L. Reciprocally, for k in $[d_K]$ such that X^{k-1} is in L, denoting $(i, j) = \psi(k)$, $\varepsilon_k Y^{i-1} Z^{j-1}$ is in L, hence $\varepsilon_k Z^{j-1}$ is in L, hence j is equal to one else this would contradict the fact that h is irreducible of degree $[K : L]$ in $L[Z]$. The set

$$\left\{\varepsilon_k X^{k-1} \,\middle|\, \exists i \in [d_L]; k = \psi^{-1}(i, 1)\right\} = \left\{Y^{i-1} \,|\, i \in [d_L]\right\}$$

is a basis of L. Hence there exist $k, k' \in \mathbb{Z}/c\mathbb{Z}$ such that $k' = k + n$ for some n in $[[K : L]]$ and that $\varepsilon_{k'} X^{k'-1}$ and $\varepsilon_k X^{k-1}$ are in L, and hence $\varepsilon_{k'}\varepsilon_k X^n$ is in L. If n were smaller than $[K : L]$, it would contradict the fact that the minimal polynomial of K over L is irreducible of degree $[K : L]$, hence $n = [K : L]$, and $\varepsilon_{k'}\varepsilon_k X^{[K:L]}$, equal to Y^{m-1} for some m in $[2; d_L]$ (m is different from one because s maps the all ones column of V to the all ones column of N), is in L. There exists $b \in \{-1; 1\}$ such that $h(Z) = Z^{[K:L]} - bY^{m-1}$.

The roots of g are hence the set

$$\left\{b \cdot \zeta_c^{k[K:L]/(m-1)} \mid k \in [d_K]; k \wedge c = 1\right\},$$

and L is a cyclotomic field of conductor ℓ and degree d_L such that $\frac{c(m-1)}{[K:L]}$ is in $\{\ell, 2\ell\}$, *i.e.* such that $[K:L]$ is in $\left\{\frac{c(m-1)}{\ell}; \frac{c(m-1)}{2\ell}\right\}$. Considering the prime decompositions of c and ℓ, $[K:L] = \frac{\varphi(c)}{\varphi(\ell)} = \frac{c}{\ell} \cdot \frac{\varphi(\mathrm{rad}(c)/\mathrm{rad}(\ell))}{\mathrm{rad}(c)/\mathrm{rad}(\ell)} \leq \frac{c}{\ell}$, so $m \leq 3$, and $[K:L]$ is in $\left\{\frac{c}{2\ell}; \frac{c}{\ell}\right\}$, or equivalently, $\frac{\mathrm{rad}(c)}{\mathrm{rad}(\ell)} \in \{1; 2\}$.

($\Longleftarrow$). If L is a cyclotomic subfield of K of conductor ℓ such that $\frac{\mathrm{rad}(c)}{\mathrm{rad}(\ell)}$ is in $\{1;2\}$, then $X^{c/\ell} = Y$, with c/ℓ equal to $[K:L]$ if $\mathrm{rad}(\ell) = \mathrm{rad}(c)$ and to $2[K:L]$ if $\mathrm{rad}(c) = 2 \cdot \mathrm{rad}(\ell)$. If $\mathrm{rad}(c) = 2 \cdot \mathrm{rad}(\ell)$, then ℓ is odd, and $X^{[K:L]} = -Y$. Finally, the minimal polynomial of K over L is $h(Z) = Z^{[K:L]} - e_{K/L}Y$, where $e_{K/L}$ is equal to -1 if $\mathrm{rad}(c) = 2 \cdot \mathrm{rad}(\ell)$ and to 1 otherwise. In the latter case, the roots of h are the primitive c-th roots of unity, and in the former one as well, considering $L' = \mathbb{Q}[Y']/(\Phi_{2\ell}(Y'))$ the 2ℓ-th cyclotomic field, such that $Y' = -Y$. There is hence a labelling $\omega_1, \ldots, \omega_{d_K}$ of the primitive c-th roots of unity such that each line i of V is equal to $\omega_i^0, \ldots, \omega_i^{d_K-1}$, and each line i of N is equal to $\mathbf{n}_i = (1, \lambda, \ldots, \lambda^{d_L-1}, \omega_i, \lambda\omega_i, \ldots, \lambda^{d_L-1}\omega_i^{[K:L]-1})$, denoting $\lambda = e_{K/L}\omega_i^{[K:L]}$.

$$\mathbf{n}_i = (e_{K/L}^{l-1}\omega_i^{(l-1)[K:L]+j-1})_{\substack{k\in[d_K], l\in[d_L], j\in[[K:L]] \cdot \\ k=(l-1)[K:L]+j}}$$

Hence, there exists a signed permutation matrix P such that $V = NP$ and $M = P$ is orthogonal, hence $\mathbb{E}\left[\|\kappa_L(A)\|^2\right] = d_K d_L \sigma^2$.

B. If $d_L = 1$. Then, as L is a subfield of K, $\mathbb{Q} \subset L$, so $L \simeq \mathbb{Q}$.

($\Longrightarrow$). Then, if $\mathbb{E}\left[\|\kappa_L(A)\|^2\right] = d_K\sigma^2$, there exists $(\varepsilon_k)_k$ in $\{-1;1\}^{d_K}$ and ψ a bijection from $[d_K]$ to $[d_K]$ such that for each k in $[d_K]$, $X^{k-1} = \varepsilon_k Z^{\psi(k)-1}$. Hence, h is of degree d_K with roots in $\left\{b \cdot \zeta_c^k \mid k \in [d_K]; k \wedge c = 1\right\}$, so h is either Φ_c or may also be Φ_{2c} if c is odd (recall that without loss of generality, we always consider c not equal to two modulo four in this article).

($\Longleftarrow$). If h is equal to Φ_c, or c is odd and h is Φ_{2c}, then there exists $(\varepsilon_k)_k$ in $\{-1;1\}$ and ψ in $\mathfrak{S}_{d_K}$ such that for each k in $[d_K]$, $X^{k-1} = \varepsilon_k Z^{\psi(k)-1}$, hence M is a signed permutation matrix and $\mathbb{E}\left[\|\kappa_L(A)^2\|\right] = d_K\sigma^2$.

□

A.3 Proof of Theorem 1

Proof. Remark that $\|\sigma(1)\|^2 = d_K = [K:L] \cdot d_L = [K:L] \cdot \|\kappa_L(1)\|^2$, so if the two forms are equal up to a scaling factor, it should be $[K:L]$. Additionally, as shown in Sections 4.1 and 4.2, if A is sampled spherically in $\kappa_{\mathbb{Q}}(A)$, with a standard deviation σ in each centred coordinate, as σ_L is trace-preserving, $\mathbb{E}[\|\sigma_L(A)\|^2] = d_K^2\sigma^2$ and $\mathbb{E}[\|\kappa_L(A)\|^2] \geq d_K d_L \sigma^2 = [K:L]^{-1}\mathbb{E}[\|\sigma_L(A)\|^2]$, with equality if and only if:

(i) L is a cyclotomic subfield of K with conductor ℓ such that $\mathrm{rad}(c)/\mathrm{rad}(\ell)$ is in $\{1;2\}$,
or (ii) L is of degree one,

as shown in Lemma 2. Hence, (i) or (ii) is a necessary condition on L. To have equality of the forms $\|\sigma_L\|^2$ and $[K:L]\|\kappa_L\|^2$, another necessary condition is that the determinant of their associated Gram matrices should be equal. For σ_L, as this embedding preserves the Euclidean norm, this determinant is Δ_K. Additionally, for $x = \sum_{i=1}^{[K:L]} x_i Z^{i-1} \in K = L[Z]/(f(Z))$ for f the minimal polynomial of L on K, $\|\kappa_L(x)\|^2 = \sum_{i=1}^{[K:L]} \sum_{\tau \in \mathrm{Hom}_{\mathbb{Q}}\, L,\mathbb{C}} \tau(x_i)\tau(\overline{x_i})$, hence the determinant of the associated Gram matrix is $\Delta_L^{[K:L]}$. A necessary condition is hence $(-1)^{d_K/2}|\Delta_K| = [K:L]^{d_K}((-1)^{d_L/2}|\Delta_L|)^{[K:L]}$, equivalent to $|\Delta_K| = [K:L]^{d_K}|\Delta_L|^{[K:L]}$, itself equivalent to (i), and to Lemma 3, (i). From Lemma 3, every a in $K = L[Z]/(Z^{[K:L]} - (-1)^{\delta}Y)$, with δ equal to $\mathrm{rad}(c)/\mathrm{rad}(\ell)$, may be written as $\sum_{i=1}^{[K:L]} a_i Z^{i-1}$ with a_i's in L. Then, the squared Euclidean norm of $\sigma_L(a)$ is

$$\begin{aligned}\|\sigma_L(a)\|^2 &= \sum_{\sigma \in \mathrm{Hom}_{\mathbb{Q}}(K,\mathbb{C})} \sigma(a)\overline{\sigma(a)} \\ &= \sum_{\substack{k \in [c] \\ k \wedge c = 1}} \sum_{i,j=1}^{[K:L]} \sum_{n,m=1}^{d_L} a_{i,n} a_{j,m} \zeta_c^{k(i-j+(n-m)[K:L])}.\end{aligned}$$

Denoting $r(i), q(i)$ the respective remainder and quotient in the Euclidean division of i by $[K:L]$,

$$\|\sigma_L(a)\|^2 = \sum_{\substack{k \in [c] \\ k \wedge c = 1}} \sum_{i,j=1}^{d_K} a_{r(i),q(i)-1} a_{r(j),q(j)+1} \zeta_c^{k(i-j)}.$$

The Euclidean norm of $\kappa_L(a)$ is

$$\|\kappa_L(a)\|^2 = \sum_{\substack{k \in [\ell] \\ k \wedge \ell = 1}} \sum_{i=1}^{[K:L]} \sum_{j,h=1}^{d_L} a_{i,j} a_{i,h} (-1)^{\delta k(j-h)} \zeta_\ell^{k(j-h)}$$

As explained in Remark 1, in the case where $2 \mid c$ and $2 \nmid \ell$, one may replace L with $L' = \mathbb{Q}[Y']/(\Phi_{2\ell}(Y'))$ of same degree, which provides the same embedding with the variable change $\ell' = 2\ell$ and $Y' = -Y$. In other cases, L' will denote L, and ℓ', ℓ, meaning that always, $\mathrm{rad}(c)$ and $\mathrm{rad}(\ell')$ have the same parity, and

$$\begin{aligned}\|\kappa_L(a)\|^2 &= \sum_{\sigma \in \mathrm{Hom}_{\mathbb{Q}}(L',\mathbb{C})} \sum_{j,k=1}^{d_{L'}} \sigma(Y'^{j-k}) \sum_{i=1}^{[K:L']} a_{i,j} a_{i,k} \\ &= \sum_{\substack{l \in [\ell'] \\ l \wedge \ell' = 1}} \sum_{j,k=1}^{d_{L'}} \zeta_\ell^{l(j-k)} \sum_{i=1}^{[K:L']} a_{i,j} a_{i,k}\end{aligned}$$

The Gram matrix corresponding to the form $\|\kappa_L\|^2$ is $(V_{K/L}P)^* V_{K/L}P$, for P a permutation matrix and $V_{K/L}$ a blockwise diagonal matrix with $[K:L]$ blocks of the Vandermonde matrix associated to L',

$$\begin{pmatrix} 1 & \zeta_{\ell'}^{k_1} & \cdots & \zeta_{\ell'}^{k_1(d_{L'}-1)} \\ 1 & \zeta_{\ell'}^{k_2} & \cdots & \zeta_{\ell'}^{k_2(d_{L'}-1)} \\ \vdots & & & \vdots \\ 1 & \zeta_{\ell'}^{k_{d_{L'}}} & \cdots & \zeta_{\ell'}^{k_{d_{L'}}(d_{L'}-1)} \end{pmatrix},$$

where $\{k_i\}_i = (\mathbb{Z}/\ell'\mathbb{Z})^\times$, of squared determinant $\Delta_{L'}$. The Gram matrix associated to the form $\|\sigma_L\|^2$ is $V_K^* V_K$, where V_K denotes the Vandermonde matrix associated to K, of squared determinant Δ_K. By definition of the Gram matrices, there is a constant $\sqrt{[K:L]}$ factor between κ_L and σ_L Euclidean norms on K if and only if for each $i, j \in [d_K]$,

- (1) for any $i, j \in [d_K]$ such that $i \neq j \mod [K:L']$,

$$\sum_{\substack{k\in[c] \\ k\wedge c=1}} \zeta_c^{k(i-j)} + \zeta_c^{-k(i-j)} = 0$$

  i.e. by reindexing and dividing by two,

$$\sum_{\substack{k\in[c] \\ k\wedge c=1}} \zeta_c^{k(i-j)} = 0,$$

  where the left-hand side is a Ramanujan sum $r_c(i-j)$ equal to [Apo76, Theorem 8.6], denoting μ the Mobius function ($\mu(n)$ is the sum of n-th primitive roots of unity),

$$\sum_{\substack{k\in[c] \\ k\wedge c=1}} \zeta_c^{k(i-j)} = \sum_{k \mid \gcd(i-j,c)} k \cdot \mu\left(\frac{c}{k}\right) = \mu\left(\frac{c}{\gcd(c,i-j)}\right) \frac{d_K}{\varphi\left(\frac{c}{\gcd(c,i-j)}\right)},$$

  where the right-hand side is Sterneck's arithmetic function in c and $i-j$. Hence, this term is null iff there is an integer $n > 1$ such that $n^2 \gcd(c, i-j)$ divides c.
- and (2), for any $i, j \in [d_K]$ such that $i = j \mod [K:L']$,

$$\begin{aligned} &\sum_{\substack{k\in[c] \\ k\wedge c=1}} \zeta_c^{k(i-j)} + \zeta_c^{-k(i-j)} = [K:L'] \sum_{\substack{h\in[\ell'] \\ h\wedge\ell'=1}} \zeta_{\ell'}^{h(q_i-q_j)} + \zeta_{\ell'}^{-h(q_i-q_j)}, \\ \iff &\sum_{\substack{k\in[c] \\ k\wedge c=1}} \zeta_c^{k(i-j)} = [K:L'] \sum_{\substack{h\in[\ell'] \\ h\wedge\ell'=1}} \zeta_{\ell'}^{h(q_i-q_j)}, \end{aligned}$$

an equality between Ramanujan sums

$$\begin{aligned}
&\iff r_c(i-j) = [K:L'] \cdot r_\ell(q_i - q_j) \\
\iff \mu\left(\frac{c}{\gcd(c, i-j)}\right) &\frac{d_K}{\varphi\left(\frac{c}{\gcd(c,i-j)}\right)} \\
&= [K:L'] \cdot \mu\left(\frac{\ell'}{\gcd(\ell', q_i - q_j)}\right) \frac{d_{L'}}{\varphi\left(\frac{\ell'}{\gcd(\ell', q_i - q_j)}\right)} \\
\iff \frac{\mu\left(\frac{c}{\gcd\left(c, [K:L'](q_i-q_j)\right)}\right)}{\varphi\left(\frac{c}{\gcd(c,[K:L'](q_i-q_j))}\right)} &= \frac{\mu\left(\frac{\ell'}{\gcd\left(\ell', q_i-q_j\right)}\right)}{\varphi\left(\frac{\ell'}{\gcd(\ell', q_i-q_j)}\right)}.
\end{aligned}$$

As we have shown that equality of the radicals of c and ℓ' was a necessary condition for the obtention of this statement, the reverse direction may now be proven using the sufficiency of meeting conditions (1) and (2). If K and L are such that $\mathrm{rad}(c) = \mathrm{rad}(\ell')$, which is equivalent to $c = [K:L']\ell'$ (Lemma 3), then

- (i) $i - j \neq 0 \mod [K:L'] \iff c \nmid (i-j)\ell'$, which denoting the prime decompositions $c = \prod_{p\in\mathcal{P}_c} p^{e_p}, \ell' = \prod_{p\in\mathcal{P}_c} p^{f_p}, i-j = \prod_{p\in\mathcal{P}_{i-j}} p^{i_p}$, means that there exists a prime p dividing c such that $e_p > i_p + f_p$ and hence $i_p \leq e_p - 2$, yielding $p^2 \gcd(c, i-j) \mid c$ and achieving condition (1);
- (ii)

$$\frac{c}{\gcd\left(c, [K:L](q_i - q_j)\right)} = \frac{[K:L'] \cdot \ell'}{[K:L'] \gcd(\ell', q_i - q_j)} = \frac{\ell'}{\gcd(\ell', q_i - q_j)},$$

fulfilling condition (2).

□

A.4 Proof of Lemma 10

Proof. First note that by linearity of expectation

$$\mathbb{E}[\log \mathrm{N}(v)] = \frac{1}{2} \sum_{\tau \in \mathrm{Hom}_{\mathbb{Q}}(K,\mathbb{C})} \mathbb{E}\left[\log(|\tau(f)|^2 + |\tau(g)|^2)\right].$$

Let τ be in $\mathrm{Hom}_{\mathbb{Q}}(K, \mathbb{C})$, which to $f(X)$ associates $f(\zeta)$. Writing $\tau(f) = a_0 + a_1\zeta + \ldots + a_{n-1}\zeta^{n-1}$,

$$\begin{cases} \mathbb{E}[|\tau(f)|^2] &= \sum_{i,j} \mathbb{E}[a_i a_j]\ \zeta^{i-j} = n\sigma^2 \\ \mathbb{E}[\tau(f)^2] &= \sum_{i,j} \mathbb{E}[a_i a_j]\ \zeta^{i+j} = R_\tau \sigma^2 \end{cases}$$

where $R_\tau = \sum_{i=0}^{n-1} \zeta^{2i} = \frac{\zeta^{2n}-1}{\zeta^2-1}$. Thus further writing $\tau(f) = U + \iota V$, for which by linearity $\mathbb{E}[U] = \mathbb{E}[V] = 0$, and as U and V are sums of Gaussian variables, they are also Gaussian, the above equations yield

$$\begin{cases} \mathbb{V}[U] + \mathbb{V}[V] = n\sigma^2, \\ \mathbb{V}[U] - \mathbb{V}[V] = \Re(R_\tau)\sigma^2, \\ 2\,\mathrm{Cov}[U, V] = \Im(R_\tau)\sigma^2. \end{cases}$$

The eigenvalues of the covariance matrix of U and V are the solutions λ to

$$\begin{aligned}(\lambda - \mathbb{V}[U])(\lambda - \mathbb{V}[V]) - \operatorname{Cov}(U,V)^2 &= 0\\ \lambda^2 - n\sigma^2\lambda + \mathbb{V}[U]\mathbb{V}[V] - \operatorname{Cov}[U,V]^2 &= 0\\ \lambda^2 - n\sigma^2\lambda + \frac{\sigma^4}{4}\left(n + \Re(R_\tau)\right)\left(n - \Re(R_\tau)\right) - \frac{\Im(R_\tau)^2\sigma^4}{2} &= 0\end{aligned}$$

and are hence

$$\lambda_1 = \frac{1}{2}n\sigma^2\left(1 - \frac{|R_\tau|}{n}\right) \qquad \lambda_2 = \frac{1}{2}n\sigma^2\left(1 + \frac{|R_\tau|}{n}\right).$$

Thus

$$|\tau(f)|^2 = U^2 + V^2 \sim \lambda_1\chi_1^2 + \lambda_2\chi_1^2.$$

Furthermore, writing $\tau(g) = U' + \iota V'$ the same computation shows that (U', V') is a two-dimensional Gaussian with the same principal axes and eigenvalues, independent from (U, V) by independence of f and g, thus

$$|\tau(f)|^2 + |\tau(g)|^2 = \lambda_1\chi_2^2 + \lambda_2\chi_2^2$$

the expected logarithm of which is

$$\frac{\lambda_1\log\lambda_1 - \lambda_2\log\lambda_2}{\lambda_1 - \lambda_2} + \psi(2) + \ln 2 - 1$$

and a lengthy computation plugging in the eigenvalues gives

$$\begin{aligned}\mathbb{E}[\log\left(|\tau(f)|^2 + |\tau(g)|^2\right)] &= \ln(n\sigma^2) + \psi(2) - \delta\left(\frac{|R_\tau|}{n}\right),\\ \delta(x) &= 1 - \frac{(1+x)\ln(1+x) - (1-x)\ln(1-x)}{2x}\end{aligned}$$

thus summing over τ gives the result, with

$$\varepsilon = \frac{1}{2}\sum_\tau \delta\left(\frac{|R_\tau|}{n}\right) = \frac{1}{2}\sum_{k\in(\mathbb{Z}/c)^\times}\delta\left(\frac{1}{n}\left|\frac{\zeta_c^{2kn} - 1}{\zeta_c^{2k} - 1}\right|\right)$$

depending only on the number field K.

Remark 6. We note [GM25, Thm. 2] states that (U, V) and (U', V') in the proof are well-approximated by centred 2-dimensional Gaussians, i.e. $\operatorname{Cov}(U, V) \approx 0$ and $\operatorname{Cov}(U', V') \approx 0$. Indeed making this approximation we recover the final result without the correction term ε.

B Comparison of Primal Attack Predictions

B.1 Proof of Theorem 2

Proof. The bijection between subfields of $K \simeq \mathbb{Q}(\zeta_c)$ and, denoting $c = \prod_{p \in \mathcal{P}_c} p^{e_p}$ the prime factor decomposition of c, subgroups of $G = \prod_{p \in \mathcal{P}_c} (\mathbb{Z}/p^{e_p}\mathbb{Z})^\times$ associates to L a subgroup H of G. We denote H_{Gal} the corresponding subgroup of $\mathrm{Gal}(K/\mathbb{Q})$, of cardinality $[K : L]$.

Denoting X the group of Dirichlet characters modulo c, one defines

$$Y = H_{\mathrm{Gal}}^{\perp} = \{\chi \in X \mid \forall h \in H_{\mathrm{Gal}},\ \chi(h) = 1\},$$

of cardinality d_L, and the root discriminant of L is provided by

$$|\Delta_L|^{\frac{1}{d_L}} = \prod_{\chi \in Y} f_\chi^{\frac{1}{d_L}},$$

where f_χ stands for the conductor of the character χ. For each character χ in X, $f_\chi = \prod_{p|c} p^{v_p(f_\chi)}$ divides c, where $p^{v_p(f_\chi)} = f_{\chi_p}$ for χ_p a character of $(\mathbb{Z}/p^{e_p}\mathbb{Z})^\times$.

$$|\Delta_L|^{\frac{1}{d_L}} = \prod_{p \in \mathcal{P}_c} \prod_{\chi \in Y} p^{v_p(f_\chi)\frac{1}{d_L}},$$

and, as $v_p(f_\chi) = |\{i \in \mathbb{N} \mid i < v_p(f_\chi)\}|$,

$$\begin{aligned}
\sum_{\chi \in Y} v_p(f_\chi) &= \sum_{\chi \in Y} |\{i \in \mathbb{N} \mid i < v_p(f_\chi)\}| \\
\frac{1}{d_L} \sum_{\chi \in Y} v_p(f_\chi) &= \frac{1}{|Y|} \sum_{i=0}^{e_p - 1} |\{\chi \in Y \mid i < v_p(f_\chi)\}| \\
&= \sum_{i=0}^{e_p - 1} 1 - \frac{|\{\chi \in Y \mid v_p(f_\chi) \leq i\}|}{|Y|} \\
\frac{1}{d_L} \sum_{p \in \mathcal{P}_c} \ln p \sum_{\chi \in Y} v_p(f_\chi) &= \sum_{p \in \mathcal{P}_c} \ln p \sum_{i=0}^{e_p - 1} 1 - \frac{|\{\chi \in Y \mid v_p(f_\chi) \leq i\}|}{|Y|}
\end{aligned}$$

$$\ln\left(\frac{|\Delta_L|^{\frac{1}{d_L}}}{|\Delta_K|^{\frac{1}{d}}}\right) = \sum_{p \in \mathcal{P}_c} \ln p \cdot A_p(X, Y),$$

where

$$A_p(X, Y) = \sum_{i \in \mathbb{N}} \frac{|\{\chi \in X \mid v_p(f(\chi)) \leq i\}|}{|X|} - \frac{|\{\chi \in Y \mid v_p(f(\chi)) \leq i\}|}{|Y|}$$

X can be decomposed as $\prod_{q\in\mathcal{P}_c} X_q$, where X_q is the set of characters of $(\mathbb{Z}/q^{e_q}\mathbb{Z})^\times$, hence,

$$\frac{|\{\chi\in X \mid v_p(f_\chi)\le i\}|}{|X|} = \frac{\left|\left\{\chi_p\in X_p \mid v_p(f_{\chi_p})\le i\right\}\right|}{|X_p|}\cdot\prod_{q\in\mathcal{P}_c;q\neq p} 1,$$

and for each i in $[e_p]$,

$$\frac{|\{\chi\in X \mid v_p(f_\chi)\le i\}|}{|X|} = \frac{\left|\widehat{(\mathbb{Z}/p^i\mathbb{Z})^\times}\right|}{\left|\widehat{(\mathbb{Z}/p^{e_p}\mathbb{Z})^\times}\right|} = \frac{\left|(\mathbb{Z}/p^i\mathbb{Z})^\times\right|}{|(\mathbb{Z}/p^{e_p}\mathbb{Z})^\times|} = \frac{\varphi(p^i)}{\varphi(p^{e_p})}.$$

Denoting, for a prime p, $Y_p = Y|_{X_p}$, there exist $y_p \mid p-1$ and $z_p < e_p$ such that $|Y_p| = y_p p^{z_p}$, and, using the convention that a $\sum_{i=0}^{-1}$ sum is null,

$$\begin{aligned}
A_p(X,Y) &= \sum_{i=z_p+1}^{e_p-1}\left(\frac{\varphi(p^i)}{\varphi(p^{e_p})}-1\right)+\sum_{i=0}^{z_p}\frac{\varphi(p^i)}{\varphi(p^{e_p})}-\frac{\gcd(y_p p^{z_p},\varphi(p^i))}{y_p p^{z_p}}\\
&= -(e_p-z_p-1)+p^{-e_p+1}\left(\frac{1}{p-1}+\sum_{i=1}^{e_p-1}p^{i-1}\right)-\frac{1}{y_p p^{z_p}}-p^{-z_p}\sum_{i=0}^{z_p-1}p^i\\
&= -(e_p-z_p-1)+\frac{1}{p-1}-\frac{1}{y_p p^{z_p}}-\frac{1-p^{-z_p}}{p-1}\\
&= -(e_p-z_p-1)-p^{-z_p}\left(\frac{1}{y_p}-\frac{1}{p-1}\right)
\end{aligned}$$

We want to show that for each prime p, $A_p(X,Y)\cdot\ln p$ is greater than $-\ln\frac{|X_p|}{|Y_p|}$ $= -\ln\left(\frac{p-1}{y_p}\right)-(e_p-1-z_p)\ln p$, *i.e.* that

$$\frac{\ln p}{p^{z_p}}\left(\frac{1}{y_p}-\frac{1}{p-1}\right)\le\ln\left(\frac{p-1}{y_p}\right).$$

For t in $]0;+\infty[$,

$$\frac{d\left(-\ln p\left(\frac{1}{t}-\frac{1}{p-1}\right)+\ln\frac{p-1}{t}\right)}{dt} = \frac{\ln p}{t^2}-\frac{1}{t},$$

is of the sign of $\ln(p)-t$, so

$$\begin{aligned}
&\min_{y_p\mid p-1}\left\{-\ln p\left(\frac{1}{y_p}-\frac{1}{p-1}\right)+\ln\frac{p-1}{y_p}\right\}\\
&\quad=\min_{y_p\in\{1;p-1\}}\left\{-\ln p\left(\frac{1}{y_p}-\frac{1}{p-1}\right)+\ln\frac{p-1}{y_p}\right\}\\
&\quad=\min\left\{\ln\left(\frac{p-1}{\left|\Delta_{\mathbb{Q}(\zeta_p)}\right|^{1/(p-1)}}\right);0\right\}=0.
\end{aligned}$$

Hence, for each prime p dividing c, $A_p(X,Y)\cdot \ln p \geq -\ln \frac{|X_p|}{|Y_p|}$, and

$$\ln\left(\frac{|\Delta_L|^{\frac{1}{d_L}}}{|\Delta_K|^{\frac{1}{d_K}}}\right) \geq \sum_{p\in\mathcal{P}_c} -\ln\frac{|X_p|}{|Y_p|} = -\ln[K:L].$$

To find which cyclotomic subfields reach this bound, we refer the reader to Lemma 3.

□

C Modelling several short vectors

Rational lattices. In the security estimates of [BN25, PT26], the square norm of the projection of a vector $\mathbf{v}\in\Lambda$ of rank N onto a projected sublattice Λ' of rank β is modelled as

$$\|\pi(\mathbf{v})\|^2 \sim \|\mathbf{v}\|^2\cdot \mathrm{Beta}\left(\frac{\beta}{2},\frac{N-\beta}{2}\right)$$

which is a consequence of heuristically assuming the subspace $W=\mathrm{span}_{\mathbb{R}}\,\Lambda'$ is uniformly random [PT26, Lem. 1]. Furthermore, if $\mathbf{v}_1,\ldots,\mathbf{v}_n$ are distinct vectors (up to sign), one may model $\|\pi(\mathbf{v}_1)\|^2,\ldots,\|\pi(\mathbf{v}_n)\|^2$ as independent beta distributions, even if *a priori* the vectors have correlated geometry [BN25, §4.2].

Module lattices. Now let K/L be an extension of cyclotomic fields and consider an $\mathcal{O}_K$-module M_K, descended canonically to an $\mathcal{O}_L$-module $M=\sigma_L(M_K)$. A short vector $\mathbf{v}\in M_K$ has rotations by roots of unity $\{\omega\mathbf{v}\}_{\omega\in\mu_K}$, which all have the same canonical norm.

Letting π_W be a projection onto a random $L_{\mathbb{R}}$-subspace W, we have that for every root of unity $\omega\in\mu_L$, the projection commutes with the action of ω,

$$\pi_W(\omega\mathbf{v}) = \omega\pi_W(\mathbf{v}) \implies \|\pi_W(\omega\mathbf{v})\| = \|\pi_W(\mathbf{v})\|$$

and thus we should *not* model $\omega\mathbf{v}$ and $\mathbf{v}$ as having independent projections – this is analogous to the caveat 'up to sign' in the case $L=\mathbb{Q}$.

However, if $\omega\in\mu_K\setminus\mu_L$ is not contained in the base field, then $\|\pi_W(\omega\mathbf{v})\|\neq\|\pi_W(\mathbf{v})\|$ in general, which suggests modelling the projections independently as before. Up to the action of μ_L, there are $|\mu_K|/|\mu_L|$ distinct rotations; assuming their projections are independent we obtain the following:

Heuristic 5. *Let M_K be an $\mathcal{O}_K$-module of rank r, let $\mathbf{v}\in M_K$, and let $M=\sigma_L(M_K)$ be the canonical descent to a subfield L, which has L-rank $s=r[K:L]$. If $W\subset\mathrm{span}_{L_{\mathbb{R}}}(M)$ is a uniformly random $L_{\mathbb{R}}$-subspace of L-rank β_L, with $1\leq\beta_L<s$, then*

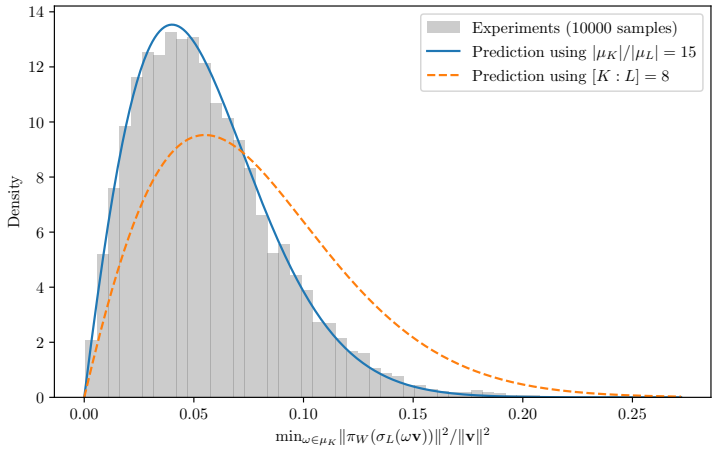
$$\min_{\omega\in\mu_K}\|\pi_W(\sigma_L(\omega\mathbf{v}))\|^2 \sim \|\mathbf{v}\|^2\cdot\min_{1\leq i\leq N} B_i$$

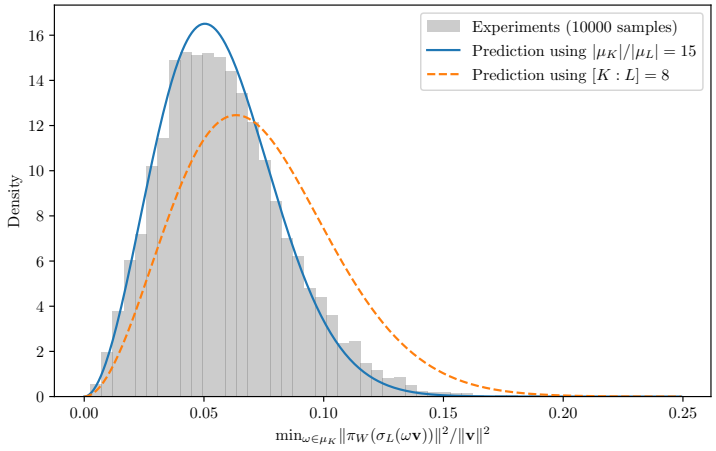
where

$$N=\frac{|\mu_K|}{|\mu_L|},\qquad B_i\overset{\text{iid}}{\sim}\mathrm{Beta}\left(\frac{d_L\beta_L}{2},\ \frac{d_L(s-\beta_L)}{2}\right).$$

One may be suspicious that $\{\omega \mathbf{v}\}_{\omega \in \mu_K}$ only spans a subspace of L-rank $[K : L]$, and so one should instead set $N = [K : L]$, as the remaining projections are L-linearly dependent to the first $[K : L]$ ones. However, we empirically test that $N = |\mu_K|/|\mu_L|$ is the correct choice in Figure 1. In any case, when K/L is a distortion-free extension then the two predictions coincide by Lemma 3.

(a) $K \simeq \mathbb{Q}(\omega_{15})$, $L \simeq \mathbb{Q}$, $r = 2$, $\beta_L = 4$





(b) $K \simeq \mathbb{Q}(\omega_{60})$, $L \simeq \mathbb{Q}(\omega_4)$, $r = 2$, $\beta_L = 3$

Figure 1: Histograms exhibiting empirical verification of Heuristic 5.