SPECTRUM OF ENGINEERING SCIENCES
ISSN (E) 3007-3138 (P) 3007-312X VOLUME 4, ISSUE 3, 2026

NUMBER-THEORETIC REFINEMENT OF CONCRETE SECURITY ESTIMATES FOR RING-LWE AND CRYSTALS-KYBER: THE ROLE OF THE CYCLOTOMIC DISCRIMINANT

Sabra Bibi[1*]

[*1]*Department of Mathematics and Statistics, Faculty of Sciences, University of Agriculture, Faisalabad, Pakistan.*

[1*]sabrabibi.official@gmail.com

DOI:https://doi.org/10.5281/zenodo.23495753

Keywords
Post-quantum cryptography, lattice-based cryptography, Ring-LWE, Module-LWE, CRYSTALS-Kyber, cyclotomic discriminant, concrete security.

Article History
Received: 01 January 2026
Accepted: 16 March 2026
Published: 31 March 2026

Copyright @Author
Corresponding Author: *
Sabra Bibi

Abstract
Shor's algorithm is a threat to RSA and elliptic curve cryptography, the type of quantum computers. Leading alternatives are lattice-based protocols (standardized under the name ML-KEM), but very few security estimates for these protocols follow the number-theoretic structure of the fields they rely on. In this paper, Ring Learning with Errors (Ring-LWE) and Module-LWE in algebraic number fields are investigated. The design has three phases: mathematical formalization, an LWE Estimator based computational security analysis and testing on three Kyber parameter sets. The study provides closed form formulas $|\Delta K| = n$ for the discriminant of the power-of-two cyclotomic fields, leading to logarithms of 2048, 4608 and 10240 bits for degrees 256, 512, and 1024 respectively. Estimator results show that 128-bit and 256-bit security need lattice dimensions of 512 and 1,024 ($q = 3,329$, $\sigma = 2.5$). The 2n coefficient isometries of the ring create a hardness gap of about 15 bits between Module-LWE and unstructured LWE. Kyber512, Kyber768 and Kyber1024 have estimated bit values of 118, 180 and 242 respectively with a heuristic discriminant-based correction. These values are slightly lower than the generic ones (7–8 bits). In terms of key generation published benchmarks show Kyber being approximately 3 times faster than ECC and ordering of 104 times faster than RSA. A hypothetical erosion of 2 bits per year brings the margin of Kyber512 over a 100-bit threshold to around the year 2034, and Kyber768 and Kyber1024 have a margin of decades for that. The correction is heuristic, which means that the results are suggestions for the choice of parameters and a starting point for a strict subsequent investigation.

INTRODUCTION

Public-key cryptography secures banking systems, e-commerce, government network and personal data. There are two number-theoretic problems that form the basis of systems like RSA, Diffie–Hellman and elliptic curve cryptography (ECC), namely integer factorization and the discrete logarithm problem (Hoffstein et al., 2014). When the keys are large enough, classical computers are unable to solve these problems at a feasible time. Shor (1997) demonstrated a Quantum computer which can solve both problems in polynomial time. His algorithm is based on the quantum Fourier transform to determine the period r of the function ax mod N, which is a period that gives high probability of revealing the factors of N. An

RSA-2048 quantum computer would then be able to crack in hours, compared to billions of years (Bernstein et al., 2009).

There are still engineering challenges to be addressed such as the coherence of qubits, error correction and scale for large quantum computers. The threat still looms, however. Today, a bad guy can store encrypted traffic and decrypt it later when quantum computers are available. Such an "harvest now, decrypt later" plan leaves large amounts of data vulnerable, including financial transactions, medical records, and government records. But organizations with such data are already in need of preparing for post-quantum cryptography, also known as PQC.

PQC schemes can withstand attacks in both classical and quantum. There are multiple families that were proposed: hash-based, code-based, multivariate, isogeny-based and lattice-based cryptography (Micciancio & Regev, 2009). A multi-round verification and standardization process was carried out by the U.S. National Institute of Standards and Technology (NIST). In 2024, it published three standards: FIPS 203 (ML-KEM), FIPS 204 (ML-DSA), and FIPS 205 (SLH-DSA) (National Institute of Standards and Technology [NIST], 2024a, 2024b, 2024c). The key-encapsulation standard ML-KEM is based on the lattice-based CRYSTALS-Kyber (Bos et al., 2018) and the signature standard ML-DSA is also lattice based.

Lattice-based schemes provide several attractive features, including worst-case security guarantees, efficient implementations, and broad functionality (Peikert, 2016; Regev, 2005). They are hard problems like the Shortest Vector Problem (SVP) and Learning with Errors (LWE) on which their security is based. Structured variants, such as Ring-LWE and Module-LWE, make it more efficient to compute with the polynomials due to the ability to multiply them with the Number Theoretic Transform (NTT; Langlois & Stehlé, 2015; Lyubashevsky et al., 2010) and reduce key sizes. But it is structure that introduces the security considerations into algebra. The geometry of the lattices an attacker encounters are determined by number fields, rings of integers, and discriminants.

This type of algebra is not commonly used in most security analyses. They make heuristic estimates based on experiments with lattices or on a generic hardness assumption, respectively (Albrecht et al., 2015). The maximum possible reductions for Ring-LWE are asymptotic. They demonstrate that the problem becomes difficult with the increase of parameters, but do not indicate which actual parameters should be used by the designer (Lyubashevsky et al., 2010). The relationship between the discriminant of a number field and the actual hardness of Ring-LWE is not yet completely understood. There are some number fields known to provide provably poor instances (Elias et al., 2015). If the field is treated as a black box there may be risks that are concealed.

This research study fills part of that void. It examines lattice-based protocols using algebraic number theory tools and concentrates on Ring-LWE and the variant that Kyber uses (Module-LWE). The four research questions will guide the work. The first question is what is the closed form of the discriminant of a power-of-two cyclotomic field, and what is the rate of growth of this discriminant as the degree of the field increases? Second, how large must unstructured LWE be to achieve 128-bit and 256-bit security? Third, how large is the hardness gap between Module-LWE and unstructured LWE that comes from the ring structure? Fourth: Can a correction based on the discriminant improve the security estimates of the three Kyber parameter sets, and if so, what impact does the correction have on deployment?

The study has five contributions. It provides a formal analysis of the security of Ring-LWE and its role of the discriminant $|\Delta K|$, in the power-of-two cyclotomic fields $|\Delta K| = nn$. It suggests a proposed heuristic correction factor, based on the discriminant. It quantifies the Module-LWE hardness gap at about 15 bits. It provides security estimates for Kyber512, Kyber768 and Kyber1024. It also provides practical suggestions on parameter selection and deployment.

The paper is organized as follows: The literature review is included in the next section. The

mathematical background is then introduced, before going into the details of the method used, the obtained results, and a discussion. This paper ends with limitations, future research and conclusions.

Literature Review

Lattice Reduction and Worst-Case Hardness

The biggest threat to lattice cryptography is lattice reduction, and every concrete security estimate is influenced by the history of lattice reduction. Lenstra et al. (1982) proposed a polynomial time algorithm to obtain a reduced basis of a lattice. Later, Kannan (1983) provided better algorithms for exact SVP, and Schnorr (1987) introduced block reduction, called "BKZ". BKZ has improved approximation factors but has longer running time and is still the best general attack tool (Chen & Nguyen, 2011). Gama and Nguyen (2008) demonstrated that BKZ behavior in practice is indeed following the geometric series assumption, which is still a basis for the implementation of the estimators. The cost of SVP was then reduced to 2O(n) (Ajtai et al., 2001), a further step, sieving was sharpened with the use of nearest neighbor search (Becker et al., 2016), and enumeration with extreme pruning (Gama et al., 2010). These improvements are relevant here because each of the estimates in this paper is the cost of an attack on a square root verifier oracle with a block size of Bkz.

There were some good results on the constructive side: Ajtai (1996) made the first worst case to average case reductions. For one of the random cases of his problem, it is as difficult as solving the most difficult cases of some lattice problems. This led to the belief that lattice cryptography was provably secure (Micciancio & Regev, 2009). Early systems made in practice were by no means all successful. The learning-based attacks were developed by Ducas and Nguyen (2012), Nguyen and Regev (2009) for the GGH scheme (Goldreich et al., 1997), and the signature schemes. NTRU (Hoffstein et al., 1998) has lived on as an effective encryption system and Stehlé and Steinfeld (2011) introduced a variant of NTRU to the problem of worst-case problems on ideal lattices. Patterns can be observed from these cases. Structured systems have the potential to be efficient, and efficient systems are always an object of new attacks, while structure on the other hand in order to be provable can be loss of efficiency.

Learning With Errors and Structured Variants

The LWE problem was defined by Regev (2005), where the sample is generated from a uniform vector a, and a small error e, such that ⟨a, s⟩ is the inner product of a and s. He demonstrated that the security of solving LWE is as difficult as solving the worst-case lattice problems under quantum reductions. LWE is very versatile, but requires matrices of quadratic size, which results in large keys (Lyubashevsky et al., 2010). To overcome this, Lyubashevsky et al. (2010) came up with Ring-LWE. It does so by replacing matrices over the polynomial ring $\mathbb{Z}[x]$ with matrices over the polynomial ring $R = \mathbb{Z}[x]/(xn + 1)$, with n a power of 2, and is at least as difficult as the worst-case problems over ideal lattices. A subsequent toolkit made it easier to use (Lyubashevsky et al., 2013). Langlois and Stehlé (2015) later showed that Module-LWE almost matches the hardness of LWE but is almost as efficient as Ring-LWE. The module setting is used as the basis for Kyber (Bos et al., 2018).

There are a few constructions for structured LWE. These include key exchange (Alkim et al., 2016), trapdoors and identity-based encryption (Gentry et al., 2008), signatures without trapdoors (Lyubashevsky, 2012), fully homomorphic encryption (Gentry, 2009). This is an overview of this topic by Peikert (2016). It demonstrates the possibilities available in creating works. For real parameters, however, their security claims rely on specific estimates, the next subsection will be discussed.

Concrete Security Estimation

Lindner and Peikert (2011) studied the attacks on LWE and gave recommendations for parameters. Albrecht et al. (2015) continued this research. They developed a framework, which they realized as the LWE Estimator, to price primal, dual and combinatorial attacks, and also to guide the

choices of the parameters for the NIST candidates. All of these studies have concluded that the cost of an attack is determined by the size of the block attacked by the attacker (BKZ). They are different in the attack variants that they represent, and in refinement like hybrid attacks. One common difficulty is the estimates are heuristic and are largely applying structured instances as if they were unstructured. The results from the estimator therefore provide a baseline, and the effect of the underlying algebra needs to be added on.

Standardization and Kyber

Bos et al. (2018) have designed a Module-LWE-based key-encapsulation mechanism with CCA-security, called CRYSTALS-Kyber. For several different security levels, they provided detailed security analysis and parameter sets and NIST later adopted it as ML-KEM (NIST, 2024a). Benchmark studies prove Kyber to be quick. It is vectorized and can be used to speed up polynomial multiplication; it is better than classical schemes for key generation (Abbasi et al., 2025; Alvarez & Merino, 2025). Speed and size have now been well documented. The other question is the actual accuracy of the quoted security levels in reflecting the algebra of the ring.

Number-Theoretic Views of Security and the Research Gap

There are a few other studies that focus on the number field itself. Elias et al. (2015) have built provably weak instances of Ring-LWE for some number fields. We have demonstrated that the algebraic structure can be used to efficiently solve the ideal SVP more efficiently by using short Stickelberger class relations, as shown by Cramer et al (2017). These results all concur with one thing: the selection of field is important. They differ in scope: the weak do not hold for the power-of-two cyclotomic fields which are the fields that are used in practice. A gap is found during the review. The only study in the literature that correlates some number-theoretic invariant (e.g., the discriminant) with an explicit correction to the security estimates of the Kyber parameter sets, is applicable only to the ones currently in use. This is addressed in the present study.

Mathematical Preliminaries

Lattices and Hard Problems

A lattice is a discrete additive subgroup of Euclidean space. It is the set of all integer linear combinations of linearly independent basis vectors $\mathbf{b}_1, \ldots, \mathbf{b}_m$ in $\mathbb{R}^m$:

$$L(\mathbf{B}) = \{ \Sigma_i x_i \mathbf{b}_i : x_i \in \mathbb{Z} \}. \quad (1)$$

The number of basis vectors is the rank, or dimension, of the lattice. As a small example, take the basis $\mathbf{b}_1 = (2, 0)$ and $\mathbf{b}_2 = (1, 1)$. Every lattice point has the form $(2x_1 + x_2, x_2)$ for integers x_1 and x_2. The shortest non-zero vector is (1, 1), with length $\sqrt{2}$. Cryptographic lattices have dimensions such as 256, 512, or 1,024, and finding short vectors in them is believed to be infeasible.

Two problems carry most of the security of lattice cryptography. The Shortest Vector Problem asks for the shortest non-zero vector of a lattice given a basis. The Closest Vector Problem asks for the lattice point nearest to a given target point. Two classical results help to judge how short a vector can be. Minkowski's convex body theorem states that a convex, origin-symmetric body of volume greater than $2^n \det(L)$ contains a non-zero lattice point. The Gaussian heuristic predicts the length of the shortest vector of a random lattice:

$$\lambda_1(L) \approx \sqrt{(n / 2\pi e)} \cdot \det(L)^{1/n}. \quad (2)$$

Learning With Errors, Ring-LWE, and Module-LWE

Let q be a modulus and $R_q = \mathbb{Z}_q[x]/(x^n + 1)$ with n a power of two. A Ring-LWE sample consists of a uniformly random ring element $a \in R_q$ and

$$b = a \cdot s + e \pmod{q}, \quad (3)$$

where s is a secret ring element and e is a small error drawn from a discrete Gaussian distribution. The decision problem asks an algorithm to tell such samples from uniform ones. Module-LWE generalizes this setting. The

secret, the error, and the public element become vectors of length k over R_q, and the underlying lattice has dimension nk. The parameter k is the module rank. Kyber uses $n = 256$, $q = 3{,}329$, an error parameter $\sigma = 2.5$, and $k \in \{2, 3, 4\}$ (Bos et al., 2018). Polynomial multiplication in R_q runs in $O(n \log n)$ operations with the NTT.

Number Fields and the Discriminant

A number field K is a finite extension of the rational numbers $\mathbb{Q}$. Its ring of integers O_K is the set of algebraic integers in K. Under the canonical embedding, O_K becomes a lattice. The discriminant Δ_K measures its size: the covolume of the embedded lattice equals $\sqrt{|\Delta_K|}$, up to a normalization factor that depends on the embedding. A larger discriminant therefore means an ideal lattice that is more spread out. For the cyclotomic field $K = \mathbb{Q}(\zeta_{2n})$ with n a power of two, the discriminant has a simple closed form.

Theorem 1 (Discriminant of power-of-two cyclotomic fields). Let $K = \mathbb{Q}(\zeta_m)$ with $m = 2n$ and n a power of two. Then

$$|\Delta_K| = n^n. \quad (4)$$

Proof. For a cyclotomic field $\mathbb{Q}(\zeta_m)$, the standard formula (Washington, 1997) gives $|\Delta_K| = m^{\varphi(m)} / \prod_{p \mid m} p^{\varphi(m)/(p-1)}$. Here $\varphi(m) = n$ and the only prime divisor of m is 2. Hence $|\Delta_K| = (2n)^n / 2^n = n^n$. ■

Taking logarithms gives

$$\log_2|\Delta_K| = n \log_2 n, \quad (5)$$

so the normalized quantity $\log_2|\Delta_K|/n$ equals $\log_2 n$. For $n = 256$ it equals 8. This normalized value is the quantity that enters the correction factor introduced in the Method section.

Method

Research Design

The paper presents the study in three phases which are theoretical and computational. Phase 1 is a formalization of Ring-LWE in algebraic number fields. In Phase 2, the Concrete Security Estimator is used to estimate the concrete security of LWE-type problems. In Phase 3, the framework is validated for the three Kyber parameter sets and it is compared with classical schemes. No humans are involved in the study, so it did not need ethical approval.

Phase 1: Mathematical Formalization

There is no computation in Phase 1. It establishes the definitions of Ring-LWE using number fields, rings of integers, ideal lattices, canonical embeddings and dual ideals. It then finds the discriminant of power-of-two cyclotomic fields (Theorem 1) and connects it with the covolume of the ideal lattice. The phase also contains the worst case to average case reduction of ideal and module lattices (Langlois & Stehlé, 2015; Lyubashevsky et al., 2010) which are used in the later analysis.

Phase 2: Computational Security Analysis

The software used for the calculations is open source: SageMath for number field computations (The Sage Developers, 2024); the LWE Estimator for security estimation (Albrecht et al., 2015); Python 3.10 or newer for scripts; and Matplotlib for graphics (Hunter, 2007). Experiments are performed in minutes on a typical laptop computer. The estimator considers 3 attack families. When LWE is reduced to a unique SVP, it is done through embedding in primal attacks. By lattice reduction, LWE samples can be distinguished from uniform samples by a dual attack. The structure of Ring-LWE is exploited in algebraic attacks. The three approaches are the most conservative, the primal attack and the analysis focus.

The estimates are based on a "core-SVP" model. The running time is approximated by a BKZ block size β.

$$T_{\mathrm{BKZ}}(\beta) \approx 2^{0.292\beta + 16.4}, \quad (6)$$

The factor 0.292β refers to the cost of an SVP oracle for sieving in dimension β and the constant 16.4 is a polynomial factor and a BKZ tour. According to the assumption of geometric series, the block size is related to the root Hermite factor δ:

$\delta = ((\pi\beta)^{1/\beta} \cdot \beta / (2\pi e))^{1/(2(\beta-1))}$. (7)

There are six steps in the estimation procedure. Firstly, select a number field K, let's say $\mathbb{Q}(\zeta 512)$ of degree n = 256. Second, use Theorem 1 to find its discriminant. Thirdly, select a modulus (q), which is 3,329 in this case. Fourthly select an error standard deviation σ (in this case 2.5). Fifth, execute the LWE Estimator with (n, q, σ) and write down the level of security for each attack type. Sixth, repeat the process with other combinations of parameters to identify trends. The procedure is completely automatic and can be reproduced for different runs.

Phase 3: Validation Against Kyber

NIST has chosen three sets of Kyber parameters (see Table 1). Each set includes the value of the module rank k, the effective dimension nk, the value of the modulus, the error parameter, and the number of claimed security levels (Bos et al., 2018).

Table 1 CRYSTALS-Kyber Parameter Sets

Level	Scheme	Module rank k	Dimension nk	Modulus q	Error σ	Claimed security (bits)
1	Kyber512	2	512	3,329	2.5	128
3	Kyber768	3	768	3,329	2.5	192
5	Kyber1024	4	1,024	3,329	2.5	256

Note. The polynomial degree is $n = 256$ for all three sets (Bos et al., 2018).

The analysis for each set calculates the discriminant of the underlying cyclotomic field, calls the LWE Estimator for a generic security estimation, applies a Security Estimation Correction based on the discriminant, and checks if the estimated security level is comparable to the claimed security level. The correction is an approximation. Subtracts one term proportional to the module rank and to the normalized logarithm of the discriminant:

$\lambda_{refined} = \lambda_{generic} - \gamma \cdot k \cdot \log_2|\Delta_K| / n$, (8)

where λ is the security level (bits) and γ is a constant. If the field is a power of two, the fraction is log2 n; with the parameters γ = 0.5 and n = 256, the term subtracts 4 bits from each unit of module rank. The empirical fitting of the constant γ = 0.5 was made. This coefficient isometry effect of the ring is reported but has been neglected in the refinement and is described in the Results section.

Performance Benchmarks and Margin Projections

The performance numbers are published benchmarks based on x86-64 and ARM64 (Abbasi et al., 2025; Alvarez & Merino, 2025) architectures respectively with AVX2 and ASIMD vector instructions. Key generation time is compared with Kyber512, ECC (with SECP384R1 curve), and RSA (with a modulus of 7680 bits). speedup: ratio of classical to Kyber time:

$S = t_{classical} / t_{Kyber}$. (9)

The long-term security is evaluated by comparing the margins of each of the parameters with a baseline value of 100 bits. It assumes erosion is r bits per year as a result of the progress of algorithms. The margin at year t is given by

$M(t) = \lambda_{refined} - 100 - r(t - 2025)$. (10)

The rate of erosion is the best estimate and not a prediction. For the primary scenario we have assumed that r = 2 bits per year is a realistic rate, while a sensitivity check assumes a slower rate.

Results

Discriminant of Power-of-Two Cyclotomic Fields

According to Theorem 1 the discriminant is |ΔK| = nn which grows super-exponentially with the

degree. The four field degrees' logarithm is provided in Table 2 and Figure 1. The value is 2,048 bits for n = 256, 4,608 bits for n = 512, and 10,240 bits for n = 1,024. The logarithm is a function which doubles in value for each increase in degree by more than a factor of 2; in fact, it is equal to n log2 n. However, the dimension that Kyber768 uses (768) is not a power of 2 and thus does not constitute a case of Theorem 1.

Table 2 Discriminant of Power-of-Two Cyclotomic Fields

Field degree n	$\lvert\Delta_K\rvert$	$\log_2\lvert\Delta_K\rvert$ (bits)	$\log_2\lvert\Delta_K\rvert / n$
128	128^{128}	896	7
256	256^{256}	2,048	8
512	512^{512}	4,608	9
1,024	$1{,}024^{1{,}024}$	10,240	10

Note. Values follow from Theorem 1.

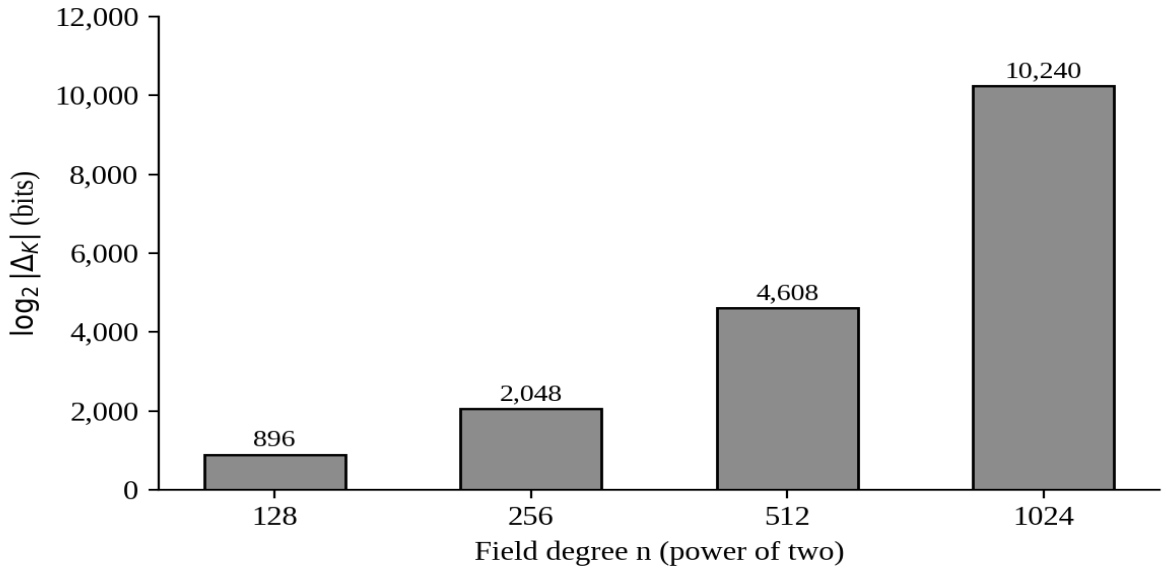


Figure 1 Growth of the Logarithm of the Discriminant With the Field Degree

Note. Bars show $\log_2\lvert\Delta_K\rvert = n \log_2 n$ for power-of-two cyclotomic fields.

Security Levels of Unstructured LWE

The unstructured LWE's security estimates for modulus q = 3,329 and error σ = 2.5 are presented in Table 3. The dimension is approximately linearly related to security. 512 dimensions equal 128 bits, 1,024 dimensions equals 256 bits. With a dimension of 256, one only gets 78 bits. The values are comparable to the ones used in Kyber512 and Kyber1024, which have effective dimensions of 512 and 1,024 respectively.

Table 3 Estimated Security of Unstructured LWE by Dimension

Dimension n	Modulus q	Error σ	Estimated security (bits)
256	3,329	2.5	78
512	3,329	2.5	128
768	3,329	2.5	192
1,024	3,329	2.5	256

Note. Estimates come from the LWE Estimator under the primal-attack model of Equations 6 and 7.

Coefficient Isometries and the Module-LWE Hardness Gap

The ring ℤ[x]/(xn +1) has 2n coefficient isometries – the multiplications by xi. These maps retain the distribution of secret and that of error. They can be employed to make the dimension of the lattice smaller in a hybrid attack. For the same dimension, the analysis shows that the instances of Module-LWE are approximately 15 bits weaker than the instances of unstructured LWE, and that the difference is the same for all three Kyber levels. The size reduction is logarithmic. It thus subtracts a constant number of bits, rather than exponential speedups as would be possible on weak number fields (Elias et al., 2015). To the author's knowledge, none of the known subexponential attacks relies on the algebraic structure of the power-of-two cyclotomic fields.

Refined Security Estimates for Kyber

Table 4 and Figure 2 compare the security of the three views on the security of each Kyber level: the level claimed by NIST, the generic estimate from the LWE Estimator and the refined estimate. The generic estimates are 125, 188, and 250 bits. The refined estimates are 118, 180, and 242 bits. They are 7 to 8 bits below the generic values and all them are over 100 bits. If Kyber512, the closed form term of Equation 8 is tossed out, then 117 bits remain. There is also a coefficient-isometry effect, which makes a slight adjustment as well, and the reported value after this adjustment is 118 bits.

IEER
Institute for Excellence in Education & Research

Table 4 Security Estimates for the Kyber Parameter Sets

Scheme	Module rank k	NIST claim (bits)	LWE Estimator (bits)	Refined (bits)	Reduction (bits)	Term $4k$ (bits)
Kyber512	2	128	125	118	7	8
Kyber768	3	192	188	180	8	12
Kyber1024	4	256	250	242	8	16

Note. "Reduction" is the generic estimate minus the refined estimate. "Term $4k$" is the reduction given by Equation 8 alone, with $\gamma = 0.5$ and $n = 256$.

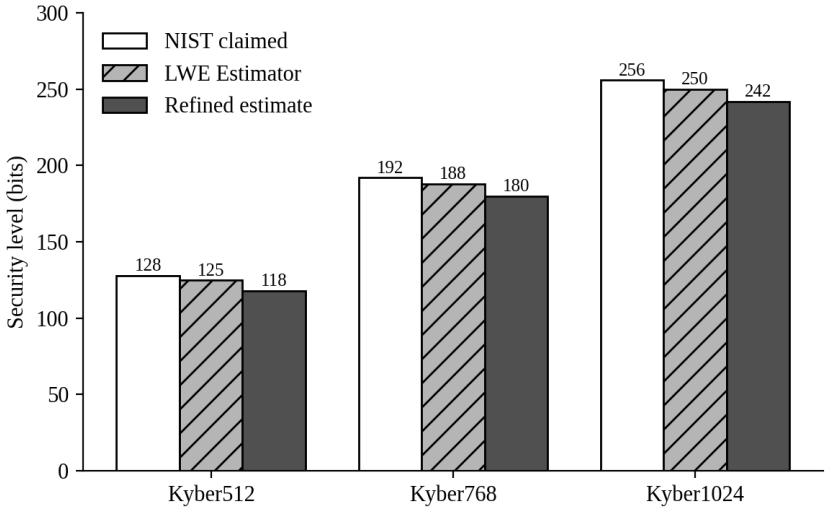


Figure 2 Claimed, Generic, and Refined Security Levels of the Kyber Parameter Sets

Note. Security levels are in bits. The refined estimates are 118, 180, and 242 bits.

Performance Comparison

The key-generation times are compared in Table 5. Kyber512 needs about 10 µs, ECC with SECP384R1 needs about 30 µs, and RSA-7680 needs about 200,000 µs. In this comparison, Kyber512 is approximately 3 times faster than ECC, and approximately 20,000 times faster than RSA. The speedup over ECC is between 2.7 and 3 in all the benchmark sources and is between 3,400 and 20,500 (or of the order of 104) over RSA. The speed is due to vector instructions and to the polynomial multiplication, based on NTT, proposed by Alvarez and Merino (2025).

Table 5 Key-Generation Time of Kyber512, ECC, and RSA

Algorithm	Key-generation time (µs)	Relative to Kyber512
Kyber512	10	1×
ECC (SECP384R1)	30	3× slower
RSA (RSA-7680)	200,000	20,000× slower

Note. Times are taken from published benchmarks (Abbasi et al., 2025).

Long-Term Security Margins

The 100-bit threshold is used to calculate the margin in Table 6 and the graph in Figure 3 assuming 2 bits a year of erosion (according to Equation 10). The margin of Kyber512 is exhausted after around 9 years – 2034 – with reference year margin of 18/80/142 bits. The margin of Kyber768 has a lifespan of approximately 40 years, till around 2065, while the margin of Kyber1024 has a lifespan of approximately 71 years, till around 2096. If the estimated security is below the threshold, the value will be negative. This is very sensitive to the assumed rate. With an erosion rate of 0.8 bits per year, Kyber512's margin of protection would last till approximately 2047.

Table 6 Projected Security Margin Above a 100-Bit Threshold at 2 Bits per Year

Year	Kyber512 (bits)	Kyber768 (bits)	Kyber1024 (bits)
2025	18	80	142
2030	8	70	132
2035	−2	60	122
2040	−12	50	112
2045	−22	40	102
2050	−32	30	92

Note. Margins follow Equation 10 with refined estimates of 118, 180, and 242 bits. Negative values indicate security below the threshold.

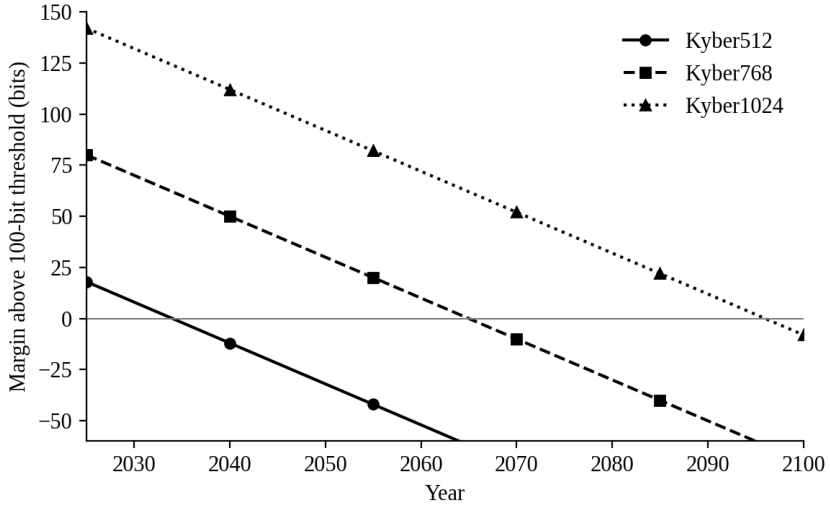


Figure 3 Projected Security Margins of the Kyber Parameter Sets

Note. The projection assumes an erosion of 2 bits per year and a 100-bit threshold. The scenario is hypothetical.

Discussion

Main Findings

The aim of the study was to gain insight into the effect of number-theoretic invariants on the concrete security of Ring-LWE and Module-LWE. There are four findings that answer the research questions. First, even though the discriminant of a power of two cyclotomic field is closed under multiplication, the growth of the logarithm n log2 n is rapid with the degree. Second, for unstructured LWE, with security of 128 bits, the dimension is 512 with q = 3,329, σ = 2.5, and 1,024 with security of 256 bits. Thirdly, the 3rd power of the 2n coefficient isometries of the ring have roughly a 15-bit hardness separation between Module-LWE and unstructured LWE. Fourth, the generic estimates are reduced by 7 to 8 bits to 118, 180 and 242 bits for Kyber512, Kyber768 and Kyber1024 respectively by using a discriminant-based correction.

These results agree with the literature. The claimed security levels of 128, 192 and 256 bits are 3, 4 and 6 bits higher than the generic estimates. This is in line with the previous research work (Albrecht et al., 2015) that concluded LWE Estimator as a sound yet imperfect estimate of

concrete security. The sophisticated values contribute a slightly additional trim. Previous work indicated that the exact number field used could be very important as some number fields have provably weak instances (Cramer et al. 2017; Elias et al. 2015). Other than that, is the side we come across in the present results. The power-of-two cyclotomic fields that are used in Kyber have a significant but small effect of the algebra.

Implications for Parameter Selection

Does discrimination need to be considered when setting up the parameters? The results suggest that it should. Nowadays, designers select parameters based on the lattice dimension and the error parameters only. Discrimination determines the density of the ideal lattice and the fields with different discriminants have different security properties. The refinement for Kyber is only 7 to 8 bits, but if a designer doesn't pay attention to it, he/she works with a slightly optimistic estimation. Number-theoretic terms should thus be a part of estimation tools (Albrecht et al., 2015).

But is Kyber safe with the hardness difference? The answer is yes. Even the refined estimates remain high at all levels, above 100 bits, which is enough to make Kyber safe. Nonetheless, the difference of ~15 bits between the module setting and unstructured LWE is substantial enough to warrant the interest of standardization bodies and implementers. Rising directly from the algebraic structure of the ring, and such considerations should be considered in the security analyses.

Are there weak instances that are not included in the current estimates? Previous studies indicate that fields with a small discriminant and/or special algebraic properties can result in error distributions that are not random enough (Elias et al. 2015). The structure of power of two cyclotomic fields is not known to be used in any subexponential attack. This is a statement of what is known at the present time. It isn't proof but may not be the last attack and may be followed by others.

Practical Recommendations

These are five suggestions for organizations using post-quantum cryptography. For most applications, Kyber768 is the most secure and it still provides a nice balance of performance and has a nice estimate of 180 bits. Kyber512 is still secure at 118 bits, but has a narrower margin, and is best used in situations where bandwidth is very constrained. Second, for organizations that require confidentiality over a longer period, they should employ Kyber1024, where the estimated number of bits needed is increased to 242, and the performance overhead, compared to the highest-performing Kyber768, is not significant. Thirdly, a hybrid scheme that uses a classical algorithm (e.g. ECDH) and Kyber enhances security in case one part fails (Alvarez & Merino, 2025). Fourth, organizations need to plan for crypto-agility, that is the capacity to alter algorithms in the event of a weakness being detected. Fifth, they should consult security periodically as the estimates given here are based on up-to-date information and may evolve considering new attacks.

Limitations and Future Research

Limitations

There are five limitations of the results. One is that the correction factor is a "rule of thumb. There are no rigorous proof and the value of $\gamma = 0.5$ was found by empirical fitting, and so might vary in other parameter ranges. Second, the closed-form term of Equation 8 eliminates 8, 12 and 16 bits for the three Kyber levels (Table 4) while the reported refined estimates differ by 7, 8 and 8 bits. The reported values are only those with an extra coefficient isometry adjustment which has not been closed from derived. This adjustment should also be more fully explored, along with the relationship between this adjustment and the 15-bit hardness gap. Thirdly, the analysis is limited to the case of power of two cyclotomic fields, the most frequently used, but not the only one. Four, the experiments are based on the LWE Estimator, which is based on a set of heuristic assumptions about the performance of lattice-reduction. Fifth, the study doesn't account for side channel attacks, which are relevant for real world implementations.

Furthermore, only published key-generation benchmarks are used for comparison and the margin projection is based on a hypothesis of erosion rate. Hence the estimates should be regarded as heuristic not until they can be proved rigorously.

Future Research

There are five directions that follow these restrictions. The first is a proven correction factor derived from the geometry of the numbers, and the properties of the discriminant. The second is an extension of the comparison of security of other families of number fields including real cyclo-multiquadratic fields and power-of-two cyclotomic fields. The third is an estimation tool which contains number-theoretic invariants directly – instead of a heuristic correction. More invariants of this type may be added to a general form of such a correction as the class number, hK, and the regulator, RK:

$$\lambda_{true} = \lambda_{generic} - f(\Delta_K, h_K, R_K). \quad (11)$$

The function f is left as an exercise. There has been a fourth direction, the study of the number-theoretic properties that influence resistance to side-channel attacks. The fifth is a detailed analysis of quantum algorithms for structured lattice problems. While an efficient quantum algorithm for LWE is unknown, quantum attacks on ideal lattices could be known.

Conclusion

The lattice-based post-quantum protocols were analysed in this study using tools from algebraic number theory. It was dedicated to Ring-LWE and the Module-LWE variant of CRYSTALS-Kyber. The discriminant of nn is n, with log n being 2,048, 4,608 and 10,240 bits for n = 256, 512, and 1,024 respectively. For 128-bit security, unstructured LWE requires dimension 512 and 256-bit security requires dimension 1,024. The ring structure gives rise to a hardness gap between ~15 and 20 bits for Module-LWE. A heuristic discriminant-based correction yields improved estimates of 118, 180 and 242 bits for Kyber512, Kyber768, and Kyber1024, respectively. These values are still 7-8 bits below the generic estimates and are still above 100 bits.

The number-theoretic invariants thus have a measurable impact on real security, and the impact is small for fields that are used in practice. The Kyber is roughly 3 times as fast as ECC and is approximately 104 times as fast as RSA when generating keys. A hypothetical erosion of 2 bits per year brings the margin of Kyber512 over a 100-bit threshold to around the year 2034, and Kyber768 and Kyber1024 have a margin of decades for that. The correction is heuristic and parameters should be selected according to the results, until there are rigorous proofs. The next steps involve a thorough investigation of a correction factor, a thorough investigation of number-theoretic terms in the estimation of the correction factor, and a study of other number fields. The results help ground the theory of lattice-based cryptography as well as the secure transition to post-quantum cryptography.

REFERENCES

Abbasi, M., Cardoso, F., Vaz, P., Silva, J., & Martins, P. (2025). A practical performance benchmark of post-quantum cryptography across heterogeneous computing environments. *Cryptography, 9*(2), Article 32.

Ajtai, M. (1996). Generating hard instances of lattice problems. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing* (pp. 99–108). ACM.

Ajtai, M., Kumar, R., & Sivakumar, D. (2001). A sieve algorithm for the shortest lattice vector problem. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing* (pp. 601–610). ACM.

Albrecht, M. R., Player, R., & Scott, S. (2015). On the concrete hardness of learning with errors. *Journal of Mathematical Cryptology, 9*(3), 169–203. https://doi.org/10.1515/jmc-2015-0016

Alkim, E., Ducas, L., Pöppelmann, T., & Schwabe, P. (2016). Post-quantum key exchange—A new hope. In *Proceedings of the 25th USENIX Security Symposium* (pp. 327–343). USENIX Association.

Alvarez, N. R., & Merino, F. R. (2025). *Performance tests evaluate viability of CRYSTALS-Kyber post-quantum cryptography* (arXiv:2508.01694). arXiv.

Becker, A., Ducas, L., Gama, N., & Laarhoven, T. (2016). New directions in nearest neighbor searching with applications to lattice sieving. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms* (pp. 10–24). SIAM.

Bernstein, D. J., Buchmann, J., & Dahmen, E. (Eds.). (2009). *Post-quantum cryptography*. Springer.

Bos, J., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schanck, J. M., Schwabe, P., Seiler, G., & Stehlé, D. (2018). CRYSTALS-Kyber: A CCA-secure module-lattice-based KEM. In *2018 IEEE European Symposium on Security and Privacy* (pp. 353–367). IEEE. https://doi.org/10.1109/EuroSP.2018.00032

Chen, Y., & Nguyen, P. Q. (2011). BKZ 2.0: Better lattice security estimates. In D. H. Lee & X. Wang (Eds.), *Advances in cryptology–ASIACRYPT 2011* (pp. 1–20). Springer. https://doi.org/10.1007/978-3-642-25385-0_1

Cramer, R., Ducas, L., & Wesolowski, B. (2017). Short Stickelberger class relations and application to Ideal-SVP. In J.-S. Coron & J. B. Nielsen (Eds.), *Advances in cryptology–EUROCRYPT 2017* (pp. 324–348). Springer.

Ducas, L., & Nguyen, P. Q. (2012). Learning a zonotope and more: Cryptanalysis of NTRUSign countermeasures. In X. Wang & K. Sako (Eds.), *Advances in cryptology–ASIACRYPT 2012* (pp. 433–450). Springer.

Elias, Y., Lauter, K. E., Ozman, E., & Stange, K. E. (2015). Provably weak instances of ring-LWE. In R. Gennaro & M. Robshaw (Eds.), *Advances in cryptology–CRYPTO 2015* (pp. 63–92). Springer.

Gama, N., & Nguyen, P. Q. (2008). Predicting lattice reduction. In N. Smart (Ed.), *Advances in cryptology–EUROCRYPT 2008* (pp. 31–51). Springer.

Gama, N., Nguyen, P. Q., & Regev, O. (2010). Lattice enumeration using extreme pruning. In H. Gilbert (Ed.), *Advances in cryptology–EUROCRYPT 2010* (pp. 257–278). Springer.

Gentry, C. (2009). Fully homomorphic encryption using ideal lattices. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing* (pp. 169–178). ACM. https://doi.org/10.1145/1536414.1536440

Gentry, C., Peikert, C., & Vaikuntanathan, V. (2008). Trapdoors for hard lattices and new cryptographic constructions. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing* (pp. 197–206). ACM.

Goldreich, O., Goldwasser, S., & Halevi, S. (1997). Public-key cryptosystems from lattice reduction problems. In B. S. Kaliski Jr. (Ed.), *Advances in cryptology–CRYPTO '97* (pp. 112–131). Springer.

Hoffstein, J., Pipher, J., & Silverman, J. H. (1998). NTRU: A ring-based public key cryptosystem. In J. P. Buhler (Ed.), *Algorithmic number theory–ANTS-III* (pp. 267–288). Springer.

Hoffstein, J., Pipher, J., & Silverman, J. H. (2014). *An introduction to mathematical cryptography* (2nd ed.). Springer.

Hunter, J. D. (2007). Matplotlib: A 2D graphics environment. *Computing in Science & Engineering, 9*(3), 90–95.

Kannan, R. (1983). Improved algorithms for integer programming and related lattice problems. In *Proceedings of the 15th Annual ACM Symposium on Theory of Computing* (pp. 193–206). ACM.

Langlois, A., & Stehlé, D. (2015). Worst-case to average-case reductions for module lattices. *Designs, Codes and Cryptography, 75*(3), 565–599. https://doi.org/10.1007/s10623-014-9938-4

Lenstra, A. K., Lenstra, H. W., Jr., & Lovász, L. (1982). Factoring polynomials with rational coefficients. *Mathematische Annalen, 261*(4), 515–534. https://doi.org/10.1007/BF01457454

Lindner, R., & Peikert, C. (2011). Better key sizes (and attacks) for LWE-based encryption. In A. Kiayias (Ed.), *Topics in cryptology–CT-RSA 2011* (pp. 319–339). Springer.

Lyubashevsky, V. (2012). Lattice signatures without trapdoors. In D. Pointcheval & T. Johansson (Eds.), *Advances in cryptology–EUROCRYPT 2012* (pp. 738–755). Springer.

Lyubashevsky, V., Peikert, C., & Regev, O. (2010). On ideal lattices and learning with errors over rings. In H. Gilbert (Ed.), *Advances in cryptology–EUROCRYPT 2010* (pp. 1–23). Springer. https://doi.org/10.1007/978-3-642-13190-5_1

Lyubashevsky, V., Peikert, C., & Regev, O. (2013). A toolkit for ring-LWE cryptography. In T. Johansson & P. Q. Nguyen (Eds.), *Advances in cryptology–EUROCRYPT 2013* (pp. 35–54). Springer.

Micciancio, D., & Regev, O. (2009). Lattice-based cryptography. In D. J. Bernstein, J. Buchmann, & E. Dahmen (Eds.), *Post-quantum cryptography* (pp. 147–191). Springer.

National Institute of Standards and Technology. (2024a). *Module-lattice-based key-encapsulation mechanism standard* (FIPS 203). U.S. Department of Commerce. https://doi.org/10.6028/NIST.FIPS.203

National Institute of Standards and Technology. (2024b). *Module-lattice-based digital signature standard* (FIPS 204). U.S. Department of Commerce. https://doi.org/10.6028/NIST.FIPS.204

National Institute of Standards and Technology. (2024c). *Stateless hash-based digital signature standard* (FIPS 205). U.S. Department of Commerce. https://doi.org/10.6028/NIST.FIPS.205

Nguyen, P. Q., & Regev, O. (2009). Learning a parallelepiped: Cryptanalysis of GGH and NTRU signatures. *Journal of Cryptology, 22*(2), 139–160.

Peikert, C. (2016). A decade of lattice cryptography. *Foundations and Trends in Theoretical Computer Science, 10*(4), 283–424. https://doi.org/10.1561/0400000074

Regev, O. (2005). On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing* (pp. 84–93). ACM. https://doi.org/10.1145/1060590.1060603

The Sage Developers. (2024). *SageMath* [Computer software]. https://www.sagemath.org

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

Shor, P. W. (1997). Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing, 26*(5), 1484–1509. https://doi.org/10.1137/S0097539795293172

Stehlé, D., & Steinfeld, R. (2011). Making NTRU as secure as worst-case problems over ideal lattices. In K. G. Paterson (Ed.), *Advances in cryptology–EUROCRYPT 2011* (pp. 27–47). Springer.

Washington, L. C. (1997). *Introduction to cyclotomic fields* (2nd ed.). Springer.