Leftover Hash Lemma(s) Over Cyclotomic Rings

Katharina Boudgoust[1] and Oleksandra Lapiha[2]

[1] CNRS, Univ Montpellier, LIRMM, France
[2] Royal Holloway, University of London, UK

katharina.boudgoust@lirmm.fr, sasha.lapiha.2021@live.rhul.ac.uk

June 9, 2025

Abstract. In this work, we propose a novel systematic approach for obtaining leftover hash lemmas (LHLs) over cyclotomic rings. Such LHLs build a fundamental tool in lattice-based cryptography, both in theoretical reductions as well as in the design of cryptographic primitives. The scattered set of prior works makes it difficult to navigate the landscape and requires a substantial effort to understand the mathematical constraints under which the LHL holds over cyclotomic rings. This is especially painful if one's given setting does not fit exactly into prior studies.

We argue that all prior approaches boil down to two different proof strategies, resulting in two main theorems. From there on, we are able to recover all previous flavours of seemingly independent LHLs as corollaries. Moreover, we showcase the power of our interpretation by providing *new* statements, covering mathematical settings not considered before. Our work further proves LHLs in the presence of *leakage* for both approaches and provides novel bounds for wide families of leakage functions.

Keywords: Lattice-Based Cryptography, Leftover Hash Lemma, Cyclotomic Rings, Leakage

Table of Contents

1 Introduction

The leftover hash lemma (LHL) is a versatile tool in cryptography, allowing to extract uniform looking randomness from a random variable which is itself not uniformly random, but has "good" properties, such as large min-entropy [ILL89,HILL99]. It can be applied in various contexts, but the one we consider in this work is its use in lattice-based cryptography. Here, the general framework can be described as follows: Let R be a cyclotomic ring and q be an integer modulus, defining the quotient ring $R_q = R/qR$. Given a vector $\mathbf{x}$ from some input distribution over R (or R_q) with guaranteed "good" properties, one obtains a uniformly looking vector over R_q by multiplying $\mathbf{x}$ from the left with a wide uniform matrix $\mathbf{A}$ over R_q. In other words, our goal is to show that $\mathbf{A}\mathbf{x} \bmod q$ (given $\mathbf{A}$) is "close" to a uniformly sampled vector $\mathbf{u}$. Most commonly, the closeness is measured using the statistical distance, but in the context of lattice-based cryptography, we also use the measure of Rényi divergence. The LHL over cyclotomic rings has been widely studied and extensively used over the past 15 years. In particular, it is of utility every time one cannot resort to computational arguments.

The Special Case of $R = \mathbb{Z}$. In lattice-based cryptography, the LHL was first used over the special ring $R = \mathbb{Z}$, i.e., the cyclotomic ring of degree 1. When q is a prime and hence $\mathbb{Z}_q$ a field, the original LHL [ILL89,HILL99] immediately holds. In more details, to apply the LHL it is enough to show that the function $\mathbf{A} \colon \mathbf{x} \mapsto \mathbf{A}\mathbf{x} \bmod q$ defines a universal hash function. That is, for a randomly chosen $\mathbf{A}$ and arbitrary inputs $\mathbf{x} \neq \mathbf{x}'$, the probability of $\mathbf{A}\mathbf{x} = \mathbf{A}\mathbf{x}'$ equals one over the size of the image. Put differently, if we fix all but one columns of $\mathbf{A}$ and $\mathbf{x} \neq \mathbf{x}'$, the remaining column is uniquely defined. This is true over $\mathbb{Z}_q$ with prime q as every non-zero element in a field is invertible. For example, the standard LHL over $\mathbb{Z}$ is used in Regev's original public-key encryption scheme [Reg05], the hardness of LWE with binary secrets [GKPV10,BLP+13], or with deterministic noise [AKPW13]. In parallel, an alternative strategy to prove the LHL was deployed in the special case where $\mathbf{x}$ follows a discrete Gaussian distribution over $\mathbb{Z}$. Such distributions arise naturally when working over lattices, hence meriting a special approach. Here, the so-called smoothing property [MR07] of discrete Gaussians is used to show that for a large enough Gaussian width, $\mathbf{A}\mathbf{x} \bmod q$ (given $\mathbf{A}$) is statistically close to a uniformly sampled vector $\mathbf{u}$. Note that the min-entropy of the Gaussian distribution scales with its width. This variant of the LHL is for instance used for proving the security of Regev's dual public-key encryption scheme [GPV08,BLR+18] and the re-randomization of LWE samples [ACPS09,GMPW20]. As far as we are aware of, [BLR+18] were the first to use the Rényi divergence as an alternative measure for closeness for the LHL over $\mathbb{Z}$.

Presence of Leakage. The standard LHL [ILL89,HILL99] can be extended to allow for additional leakage on the vector $\mathbf{x}$, as proven in [DORS08]. This is commonly referred to as *generalised* LHL. More precisely, the goal is now to show that $\mathbf{A}\mathbf{x} \bmod q$, given $\mathbf{A}$ as well as some additional leakage y on $\mathbf{x}$, is close to a uniformly sampled vector $\mathbf{u}$. It requires proving that the distribution of $\mathbf{x}$ given the leakage y still has "good" properties. Concretely, it is sufficient to show that the average conditional min-entropy of $\mathbf{x}$ given y is high. The generalised LHL can easily be applied to the same use cases over $\mathbb{Z}$ described above. It served for instance to demonstrate the hardness of LWE with entropic secrets [BD20] or for arguing the circuit privacy of FHE [BdPMW16].

The General Case. Over time, the lattice-based community moved to cyclotomic rings of larger degrees as they allow for significantly more compact cryptographic schemes as well as faster computations, e.g. [SSTX09,LPR10,LS15]. However, when using cyclotomic rings of degrees larger than 1, one cannot apply the LHL in the same way as before. In general, R_q is not a field, even if q is prime.[3] Hence, not every non-zero ring element is invertible. This observation marked the start of a long sequence of works, many claiming a novel LHL over cyclotomic rings. To the best of our knowledge, the first LHL over (not necessarily cyclotomic) rings is due to Micciancio [Mic07]. It was originally restricted to cyclic rings of the form $\mathbb{Z}[X]/\langle X^N - 1\rangle$ for any integer N and only applied to matrices $\mathbf{A}$ with one row. It was subsequently generalised to cyclotomic rings [SSTX09] and matrices of multiple rows [BJRW23]. The latter work was also the first to use the Rényi divergence measure in the general case. In [KY16], limiting the setting to power-of-2 cyclotomic rings with a two-splitting prime modulus[4] allowed for tighter parameters. All quoted works only consider the uniform distribution over elements of small norm as the input distribution. Later, [LW20] claimed a different proof approach and obtained a more general result. Their LHL applied to the ring of integers of any number field and arbitrary hash input distributions, as long as they had a small enough collision probability (or large enough min-entropy). Note that all the results listed until now do not allow for any leakage. The work of [LW20] was generalised to a specific leakage function (namely, leaking $\mathbf{x} + \mathbf{e}$, where $\mathbf{e}$ follows a continuous Gaussian distribution) in [LWZW24]. In a recent follow-up work [JLWG25], the LHL of [LWZW24] was generalised to arbitrary leakage functions. In [BI22], another version of the LHL is proven tailored to the setting of power-of-2 cyclotomics with power-of-2 moduli and noisy linear leakage.

In parallel to generalising the original LHL proof to the structured setting, another line of work [SS11,LPR13,RSW18] aimed at extending the smoothing-based argument for discrete Gaussians to higher degree cyclotomics. Even though the high level approach stays the same, the main challenge lies in computing the smoothing parameter (to be defined later) of a kernel lattice corresponding to the matrix $\mathbf{A}$. The computations are significantly easier for integer matrices compared to matrices over higher-degree cyclotomic rings. Additionally, there are two ways of sampling a discrete Gaussian distribution over R, either via the coefficient embedding or the canonical embedding, and the smoothing analysis depends on this choice. Whereas the generalised LHL can handle arbitrary leakage, allowing for leakage in the smoothing-based approach is less obvious. So far, it had only been addressed for specific leakage functions [DGKS21,SSE$^+$24].

Ideally, one aims for an LHL over cyclotomic rings as versatile as the standard LHL over finite fields. Moreover, the application of such an LHL should not require the user to be an expert on cyclotomic rings. Unfortunately, the scattered landscape of LHLs over cyclotomic rings, each of them having a different setting and parameter restrictions, makes a black-box use very hard. This brings us to the following research questions we asked ourselves in this work:

> *Is it possible to 1) generalise the seemingly independent flavours of the LHL over cyclotomic rings into few single results, which 2) allow an easy adaptation to new mathematical settings and 3) new leakage functions?*

Our Results. We positively answer the above questions and show that the large set of LHLs over cyclotomic rings can be reduced to two main theorems. The first approach leads us to the LHL *via counting ideals*, and the second one to the LHL *via smoothing*. The previous results can then be expressed as corollaries of one of the two theorems. We further generalise both approaches to the leakage context, providing novel *generalised* LHL's over cyclotomic rings. We highlight that the counting approach applies to a large class of input distributions (including discrete Gaussians), whereas the smoothing approach is so far limited to discrete Gaussian distributions. Still, the latter approach has its benefits as it allows for parameter configurations not covered by the counting proof technique, as well as a tailored analysis in the case of linear leakage. We give

[3] Requiring R_q to be a field is a very restrictive condition. For power-of-2 cyclotomic rings, predominant in practise, there exists no integer q meeting the condition.

[4] That is, the ring R_q is the product of two finite fields.

more details on this in the technical overview. With this new perspective at hand, we are able to extend many of the previous results to broader settings, meaning that we can use different distributions and splitting behaviours, avoiding prior limitations.

We would like to emphasise once more the fundamental nature of the LHL in cryptography, justifying its study on its own. Nonetheless, we see some concrete applications of our new results, especially when considering practical cryptosystems. All NIST lattice standards use fully-splitting rings. For such rings, besides bounded uniform and discrete Gaussian input distributions, no LHL was known prior to our work. Due to implementation considerations, many prefer other distributions, e.g. the central binomial distribution as used in Kyber [SAB+22]. Our new LHL presented in Corollary 7 is the first to apply to fully-splitting rings and such distributions used in practice.

Our Scope. Throughout this work, we focus on the LHL over cyclotomic rings that proves $\mathbf{A}\mathbf{x} \bmod q$ (given $\mathbf{A}$ and potentially some leakage on $\mathbf{x}$) is close to uniform, where $\mathbf{A}$ comes from the uniform distribution. We thus do not look at results where $\mathbf{A}$ is kept secret, like [CPS+20], or where both $\mathbf{A}$ and $\mathbf{x}$ are Gaussians and are shown to be again close to some Gaussian distribution, like [AGHS13,LSS14]. Further, we do not study computational arguments, where the closeness of $(\mathbf{A}, \mathbf{A}\mathbf{x})$ to uniform is guaranteed by relying on hardness assumptions related to structured lattices, such as Module LWE [LS15]. We concentrate on cyclotomic fields, the ones most used in lattice-based cryptography, but highlight if results generalise to more general number fields.

1.1 Technical Overview

Let us give some more details regarding our key insights. Our overall goal is to give concrete constraints under which $(\mathbf{A}, \mathbf{A}\mathbf{x})$ is close[5] to $(\mathbf{A}, \mathbf{u})$, where the $n \times m$ matrix $\mathbf{A}$ and the vector $\mathbf{u}$ are sampled uniformly at random over R_q for a cyclotomic ring R and a modulus q. The hash input vector $\mathbf{x}$ is sampled from some distribution $\mathcal{P}$ over R_q^m (or over R^m).

Via Counting Ideals. The first approach makes use of the well-known observation that a small collision probability of $(\mathbf{A}, \mathbf{A}\mathbf{x})$ implies closeness to uniform. Upper bounding the first boils down to bounding the probability $\Pr[\mathbf{A}(\mathbf{x} - \mathbf{x}') = \mathbf{0}]$ for uniform random $\mathbf{A}$ and independently sampled $\mathbf{x}, \mathbf{x}' \leftarrow \mathcal{P}$. A basic result of algebra (Lemma 10) states that for fixed $\mathbf{x}, \mathbf{x}'$ and uniform random $\mathbf{A}$, the expression $\mathbf{A}(\mathbf{x} - \mathbf{x}')$ is distributed uniformly at random over the ideal $\mathcal{I}_{\mathbf{x},\mathbf{x}'} \subseteq R_q$ generated by $\mathbf{x} - \mathbf{x}'$, that is $\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \langle x_1 - x_1', \ldots, x_m - x_m' \rangle$. Using that the uniform distribution over some ideal $\mathcal{I}$ is proportional to its norm $N(\mathcal{I})$ and conditioning (and hence *counting*) over all ideals of R_q, we obtain our first LHL.

Theorem 1 (Counting Ideals, Informal). *Let $\mathbf{A} \leftarrow \mathcal{U}(R_q^{n \times m})$ and $\mathbf{x} \leftarrow \mathcal{P}$ for some distribution over R_q^m. The distance of $(\mathbf{A}, \mathbf{A}\mathbf{x})$ to the uniform distribution scales with*

$$\sum_{\mathcal{I} \neq R_q} N(\mathcal{I}) \cdot \Pr_{\mathbf{x},\mathbf{x}' \leftarrow \mathcal{P}}[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I}]. \tag{1}$$

In other words, whether a distribution $\mathcal{P}$ is good for randomness extraction depends on the behaviour of $\mathcal{I}_{\mathbf{x},\mathbf{x}'}$ for independently sampled $\mathbf{x}, \mathbf{x}' \leftarrow \mathcal{P}$. The formal statement can be found in Theorem 4. We want to emphasise that the use of the ideal $\mathcal{I}_{\mathbf{x},\mathbf{x}'}$ itself is not novel, but that our contribution lies in realising that the above expression builds the common starting point of various seemingly independent LHLs, allowing us to derive previous and new flavours.

[5] For simplicity, we will not distinguish between statistical and Rényi closeness here.

By observing that $\Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I}] \leq \Pr[\forall i : x_i - x_i' \in \mathcal{I}] = \mathrm{Coll}(\mathcal{P} \bmod \mathcal{I})$ and $\mathrm{Coll}(\mathcal{P} \bmod \mathcal{I}) \leq 2^{-\mathrm{H}_\infty(\mathcal{P} \bmod \mathcal{I})}$, we replace the dependency on the rather algebraic object $\mathcal{I}_{\mathbf{x},\mathbf{x}'}$ by more well-studied objects like collision probability and min-entropy. Here, $\mathcal{P} \bmod \mathcal{I}$ denotes the distribution of sampling $\mathbf{x} \leftarrow \mathcal{P}$ and taking each of its coefficients modulo $\mathcal{I}$. This leads to the LHL presented in [LW20] and formulated in Corollary 3. Note that described proof techniques are very general and the results so far apply to *any* finite ring.

Next, we use known upper bounds on the collision probability and/or min-entropy for specific distributions over cyclotomic rings to give more concrete formulas for the expression in Equation (1). Previously, this has been done for bounded uniform distributions in [Mic07,SSTX09,BJRW23], as covered by Corollary 4. By using recent min-entropy bounds of [JLWG25], we provide a tailored formula (implicit in [JLWG25]) for discrete Gaussian distributions in Corollary 5.

By trivially upper bounding the norm $N(\mathcal{I})$ by the size of R_q, it is enough to compute the probability $\sum_{\mathcal{I} \neq R_q} \Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I}] = \Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} \neq R_q]$. If we are in the low-splitting regime[6] and $\mathcal{P}$ provides short-norm elements, we can use [LS18] to bound the latter probability by $\mathrm{Coll}(\mathcal{P})$, leading to Corollary 6. More precisely, [LS18] guarantees that as long as $\mathbf{x} \neq \mathbf{x}'$, the ideal $\mathcal{I}_{\mathbf{x},\mathbf{x}'}$ contains a unit and hence generates the full ring R_q. Thus, we get $\mathcal{I}_{\mathbf{x},\mathbf{x}'} \neq R_q$ only when $\mathbf{x}$ and $\mathbf{x}'$ are equal. In general, it is easier to bound $\mathrm{Coll}(\mathcal{P})$, rather than $\mathrm{Coll}(\mathcal{P} \bmod \mathcal{I})$ for any possible ideal $\mathcal{I}$. Hence, Corollary 6 is easier to apply than Corollary 3, at the cost of imposing a low-splitting regime. A similar reasoning was done in [KY16], but restricted to power-of-2 cyclotomics and the two-splitting regime. Our Corollary 6 generalises it to any cyclotomic ring and any low-splitting regime.

Last but not least, with Corollary 7 we prove a new LHL based on the well-spreadness property of $\mathcal{P}$. The latter has been first formalised in [AAB+24]. Informally, it bounds the probability of an element sampled from $\mathcal{P}$ hitting a specific element in each CRT slot of R_q. Even though in general well-spreadness of $\mathcal{P}$ leads to a looser statement compared to the collision probability of $\mathcal{P}$, it can be used even when R is splitting in many factors modulo q. In other words, Corollary 7 applies in a high-splitting regime, whereas Corollary 6 is limited to a low-splitting ring. To the best of our knowledge, this is the first LHL applying in the high-splitting regime for input distributions other than bounded uniform and discrete Gaussians.

In Section 5, we provide bounds for collision probability and well-spreadness of various distributions studied in the lattice literature. In particular, we cover distributions beyond the bounded uniform and discrete Gaussian case. For example, we look at the centered binomial distribution (as in the NIST standard Kyber [SAB+22]) or the ternary distribution with fixed Hamming weight.

The counting ideals approach can be generalised to the setting of leakage on $\mathbf{x}$. More concretely, our goal is to show that $(\mathbf{A}, \mathbf{A}\mathbf{x}, y)$ is close to $(\mathbf{A}, \mathbf{u}, y)$, where $\mathbf{A}, \mathbf{u}$ are uniform over R_q, $\mathbf{x} \leftarrow \mathcal{P}$ and $y \leftarrow \mathcal{Q}$. Here, $\mathcal{Q}$ depends on $\mathbf{x}$. Again, there exists some (maybe less well-known) relation between closeness-to-uniform and the *average conditional* collision probability $\mathbb{E}_{y \leftarrow \mathcal{Q}}[\mathrm{Coll}((\mathbf{A}, \mathbf{A}\mathbf{x}) \mid y)]$.[7] We compute this expected value in a similar way as in the non-leakage case. In particular, we now have to bound the probability $\mathbb{E}_{y \leftarrow \mathcal{Q}}[\Pr[\mathbf{A}(\mathbf{x} - \mathbf{x}') = \mathbf{0} \mid y]]$. Applying Lemma 10, conditioning over all ideals of R_q and by linearity of the expectation, we obtain the following informal leakage analogue of Theorem 1.

Theorem 2 (Counting Ideals with Leakage, Informal). *Let $\mathbf{A} \leftarrow \mathcal{U}(R_q^{n \times m})$ and $\mathbf{x} \leftarrow \mathcal{P}$ for some distribution over R_q^m. Further, let $\mathcal{Q}$ be a distribution depending on $\mathbf{x}$. The distance of*

[6] We say that the ring R is low-splitting modulo q, if R_q is the product of few fields.

[7] There exist different definitions of average conditional collision probability in the literature. Our choice fits well with the use of the Rényi divergence measure.

$(\mathbf{A}, \mathbf{Ax})$*, given* $y \leftarrow \mathcal{Q}$*, to the uniform distribution scales with*

$$\sum_{\mathcal{I} \neq R_q} N(\mathcal{I}) \cdot \mathbb{E}_{y \leftarrow \mathcal{Q}}[\Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I} \mid y]]. \tag{2}$$

Hence, whether a distribution $\mathcal{P}$ is good for randomness extraction depends on the expected value of $\Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I}]$ under leakage y. The formal statement can be found in Theorem 5. Note that when $\mathcal{Q}$ is independent of $\mathbf{x}$, we recover Equation (1) from Theorem 1. Again, this result builds the starting point to derive previous and new flavours of the LHL.

Let $\widetilde{\mathrm{Coll}}$ and $\widetilde{\mathrm{H}_\infty}$ denote the average conditional collision probability and min-entropy defined in the main body. Using that $\mathbb{E}_{y \leftarrow \mathcal{Q}} \Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I} \mid y] \leq \widetilde{\mathrm{Coll}}(\mathcal{P} \bmod \mathcal{I} \mid \mathcal{Q}) \leq 2^{-\widetilde{\mathrm{H}_\infty}(\mathcal{P} \bmod \mathcal{I} | \mathcal{Q})}$, we recover Corollary 3, which can be seen as a generalisation of [LWZW24] to arbitrary leakage and of [BI22] to arbitrary leakage, cyclotomic rings and moduli. It was recently proven [JLWG25].

For the low-splitting regime, we observe that the invertibility properties proven in [LS18] do not depend on the exact distribution, and hence are valid even in the presence of leakage. Therefore, we obtain an expression that depends only on $\widetilde{\mathrm{Coll}}(\mathcal{P} \mid \mathcal{Q})$. This results in a novel LHL in the low-splitting regime, Corollary 6. It can be seen as a generalisation of [LWZW20] from the extremely limiting setup requiring R_q to be a field to any low-splitting regime.

Note that we do not propose a leakage variant of Corollary 7, as well-spreadness does depend on the distribution. Such a generalisation requires a notion of well-spreadness under leakage, which we leave for future research.

Via Smoothing. The second approach focuses on the special case where the hash input $\mathbf{x}$ comes from a discrete Gaussian distribution over R^m of width σ, denoted $\mathcal{D}_{R^m,\sigma}$.[8] As Gaussians arise very naturally in lattice-based cryptography, a dedicated line of research has evolved. It starts by observing that the closeness of $(\mathbf{A}, \mathbf{Ax})$ to uniform can be expressed as the expected value over $\mathbf{A}$ of the closeness of $\mathbf{Ax}$ to uniform. For example for the statistical distance SD, it holds $\mathrm{SD}((\mathbf{A}, \mathbf{Ax}), (\mathbf{A}, \mathbf{u})) = \mathbb{E}_{\mathbf{A}}[\mathrm{SD}(\mathbf{Ax}, \mathbf{u})]$. Then, for a fixed $n \times m$ matrix $\mathbf{A}$, the map $\mathbf{A} \colon \mathbf{x} \in R^m \mapsto \mathbf{Ax} \bmod q \in R_q^n$ induces an isomorphism between $R^m / \Lambda_q^\perp(\mathbf{A})$ and $\mathrm{im}(\mathbf{A})$, where $\Lambda_q^\perp(\mathbf{A}) = \{\mathbf{y} \in R^m \mid \mathbf{Ay} = \mathbf{0} \bmod q\}$ is the kernel of the map. For simplicity, let us assume for now that $\mathrm{im}(\mathbf{A}) = R_q^n$. Then, $\mathbf{Ax} \bmod q$ is close to uniform over R_q^n if and only if $\mathbf{x} \bmod \Lambda_q^\perp(\mathbf{A})$ is close to uniform over $R^m / \Lambda_q^\perp(\mathbf{A})$. At this point, we use a well-known observation on discrete Gaussian distributions made in [GPV08]: closeness to the uniform distribution is guaranteed as long as the Gaussian width of $\mathbf{x}$ lies above the smoothing parameter η of $\Lambda_q^\perp(\mathbf{A})$. Informally, the smoothing parameter [MR07] provides a threshold of the Gaussian width, above which we have a much better understanding how discrete Gaussians distributions behave. Overall, this gives us the following *smoothing-based* LHL.

Theorem 3 (Smoothing, Informal). *Let* $\mathbf{A} \leftarrow \mathcal{U}(R_q^{n \times m})$ *and* $\mathbf{x} \leftarrow \mathcal{D}_{R^m,\sigma}$*. The distribution* $(\mathbf{A}, \mathbf{Ax})$ *is close to the uniform distribution if both* $\mathrm{im}(\mathbf{A}) = R_q^n$ *and* $\sigma \geq \eta(\Lambda_q^\perp(\mathbf{A}))$ *with overwhelming probability over the choice of* $\mathbf{A}$*.*

The formal statement can be found in Theorem 6. The existing literature then deviates on *how* an upper bound the smoothing parameter of the kernel lattice is obtained. The works of [SS11,RSW18] derive a lower bound on the first minimum of the dual lattice. To do that they use a union bound over all possible vectors within a ball, paired with a volume argument for the ball and a set of ideal lattices. In contrast [LPR13] proceed by the definition of the smoothing parameter and directly compute the expected Gaussian weight of the dual lattice and then apply

[8] Our results apply to non-spherical discrete Gaussians with non-zero center, but for simplicity we focus on spherical and zero-centered in the technical overview.

Markov's inequality. Interestingly, even though the proofs differ substantially, both use the same technical result (Lemma 10) over the ideal $\mathcal{I}_{\mathbf{z}}$ for some vector $\mathbf{z}$ over R_q, exactly the same result which was a key observation in the counting approach.

Note that for a matrix $\mathbf{A}$ over the ring R_q, interpreting $\Lambda_q^{\perp}(\mathbf{A})$ as a lattice, requires to specify which of the two possible embeddings is used. Concretely, [SS11] opted for the coefficient embedding, and [LPR13,RSW18] state their results in the canonical embedding. Whereas the latter applies to any cyclotomic ring, the coefficient version is limited to power-of-2 cyclotomics. This constraint comes from the fact that for the smoothing analysis, [SS11] needs to express the dual of $\Lambda_q^{\perp}(\mathbf{A})$ as the lattice $\Lambda_q(\mathbf{A}')$ of some related matrix $\mathbf{A}'$ over R_q. We show with Corollary 8 that this limitation is not necessary. Our key observation is that once we have the smoothing-based LHL (including the smoothing analysis for $\Lambda_q^{\perp}(\mathbf{A})$) in one embedding, we can switch to the other embedding (without the need of re-doing the smoothing analysis of $\Lambda_q^{\perp}(\mathbf{A})$ for the other embedding). Depending in which direction we go, the incurred loss is determined by the smallest or largest singular value of the Vandermonde matrix of the given cyclotomic ring. In the case of power-of-2 cyclotomics, we exactly recover [SS11] from the canonical embedding statement. To the best of our knowledge, this is the first smoothing-based LHL in the coefficient embedding that applies to any cyclotomic ring.

Regarding the image condition, all prior results enforce $\operatorname{im}(\mathbf{A}) = R_q^n$ by either prepending the identity matrix [LPR13], recalled in Lemma 24, or by limiting to invertible entries [SS11,RSW18], restated in Lemma 22+23. We provide new LHL versions for both approaches in the more natural case when the matrix has to be truly uniform (cf. Lemma 25+26). This comes at the expense of either requiring a low-splitting regime or non-constant number of columns m of $\mathbf{A}$, similarly to the counting approach.

Our work concludes with a study of the smoothing-based approach when linear leakage on the input vector $\mathbf{x}$ is given. We distinguish between *exact* and *noisy* linear leakage. The noisy leakage is of the form $(\mathbf{E}, \mathbf{E}\mathbf{x} + \mathbf{y})$, where the matrix $\mathbf{E}$ is given, but $\mathbf{y}$ stays hidden and comes from a different discrete Gaussian distribution. We assume $\mathbf{E}$, $\mathbf{x}$ and $\mathbf{y}$ to all have short norm, and do not perform the computation modulo q. By using recent results from the literature [ENP24], we observe that the distribution of $\mathbf{x}$ (originally coming from a zero-centred discrete Gaussian distribution) given the leakage $(\mathbf{E}, \mathbf{E}\mathbf{x} + \mathbf{y})$ still follows a discrete Gaussian distribution over the ring. The width and centre of the new distribution depend on the leakage, but $\mathbf{x}$ itself is sampled independently. In other words, we can apply the normal smoothing-based LHL (Theorem 3), where $\mathbf{x}$ comes from the discrete Gaussian distribution defined by the leakage, leading to our noisy leakage result in Theorem 8. We note that the first type of leakage function considered in [DGKS21] can be seen as a special case, where $\mathbf{E}$ is the identity matrix. Note that the two other types of leakage in [DGKS21] are specific to the side-channel attack setting, and we do not cover them in this work.

In the exact leakage context, the hint is of the form $(\mathbf{E}, \mathbf{E}\mathbf{x})$, where $\mathbf{x}$ comes from a discrete zero-centred Gaussian distribution. Again, there will be no modulo q. The hint can be seen as a constraint, conditioning $\mathbf{x}$ to lie in a specific coset of the kernel lattice $\Lambda^{\perp}(\mathbf{E})$. Put differently, $\mathbf{x}$ follows a discrete Gaussian distribution over the lattice $\Lambda^{\perp}(\mathbf{E})$, and we want to show that $\mathbf{x} \bmod \Lambda^{\perp}(\mathbf{E}) \cap \Lambda_q^{\perp}(\mathbf{A})$ is close to uniform. As in the non-leakage context, we can use a smoothing argument once the Gaussian width of $\mathbf{x}$ lies above the smoothing parameter of the new lattice $\Lambda^{\perp}(\mathbf{E}) \cap \Lambda_q^{\perp}(\mathbf{A})$. The remainder of the proof gives a bound on the latter, following the approach of [SS11], and leading to Theorem 7. We remark that our proof technique is inspired by the specific leakage context in [SSE^{+}24]. However, they only applied it in a noisy leakage context, and did not detail its potential use for the more powerful *exact* leakage. In fact, we provide the first meaningful bound for the LHL with exact leakage as previously its effect was bounded by its bit-size. In practice, this means that in prior work leaking even one ring element would increase the statistical distance to uniform by a factor exponential in the ring degree (see Lemma 9).

Counting vs. Smoothing. The attentive reader might be wondering why we keep the smoothing approach all along, given that the counting approach already applies to discrete Gaussians, cf. Corollary 5. We detail two aspects where the smoothing approach beats the counting approach. First, as remarked in [LPR13], if the matrix $\mathbf{A}$ is sampled from the uniform distribution, then both the smoothing and counting LHL cannot allow for both the number m of columns of $\mathbf{A}$ *and* the degree δ of the splitting to be constant. In other words, we cannot *simultaneously* have constant matrix dimensions and be in the high-splitting regime. However, in the smoothing approach, we can circumvent this limitation, for instance by prepending the identity matrix to $\mathbf{A}$. An additional advantage of the smoothing approach over the counting approach is that for exact linear leakage, we obtain tighter results and hence allow for more information to be leaked.

Open Problems. The last missing step to systemise the study of leftover hash lemmas over cyclotomic rings, would be to find a general theorem encompassing both the counting and the smoothing approach. Moreover, tighter min-entropy bounds modulo ideals and under specific leakage distributions would yield better LHL results. Extending the smoothing-based approach to distributions other than discrete Gaussian and to larger classes of leakage, beyond linear leakage, could be an interesting future direction. Finally, a curious algebraic problem would be to analyse the probability of $\langle x_1 - x'_1, \ldots, x_m - x'_m \rangle = \mathcal{I}$ in a tighter way than resorting to $\Pr[\forall i : (x_i - x'_i) \in \mathcal{I}]$.

2 Preliminaries

We denote matrices in uppercase boldface letters $\mathbf{A}$ and vectors in lowercase $\mathbf{x}$. We write $[\mathbf{A} \mid \mathbf{B}]$ for concatenating matrices horizontally and $[\mathbf{A}||\mathbf{B}]$ for stacking them vertically. For $\mathbf{A} \in \mathbb{R}^{n \times m}$ we write $\|\mathbf{A}\|_F = \sqrt{\sum_{i,j} a_{ij}^2}$ to signify the Frobenius norm of the matrix. The $\mathfrak{s}_1(\mathbf{A})$ and $\mathfrak{s}_m(\mathbf{A})$ denote the largest and the smallest singular values of $\mathbf{A}$ respectively. It holds that $\mathfrak{s}_1(\mathbf{A}) \leq \|\mathbf{A}\|_F$. We denote the Euclidean ℓ_2 norm as $\|\mathbf{x}\|$ and the infinity norm as $\|\mathbf{x}\|_\infty$.

2.1 Number Theory

A number field $K = \mathbb{Q}(\zeta)$ of degree N is a finite field extension of the rationals $\mathbb{Q}$ obtained by adjoining an algebraic number ζ. We denote its ring of integers by R. We call K a ν-th cyclotomic number field if ζ is a ν-th primitive root of unity. Its degree is given by $N = \varphi(\nu)$, where φ is Euler's totient function.

We can identify $K = \mathbb{Q}[X]/\langle \Phi(X) \rangle$, where $\Phi(X)$ is the minimal polynomial of ζ. Every element $x \in K$ can then by written with respect to the basis $\left\{1, \zeta, \ldots, \zeta^{N-1}\right\}$, thus $x = \sum_{i=0}^{N-1} x_i \zeta^i$ with $x_i \in \mathbb{Q}$. The isomorphism $\tau \colon K \to \mathbb{Q}^N$ which maps x to its coefficient vector $\tau(x) = (x_0, \ldots, x_{N-1})^T$ is called the coefficient embedding. By associating the norm of an element x in K with the norm of its corresponding $\tau(x) \in \mathbb{Q}^N$, it is possible to equip K with a geometry. For a positive integer η, we define $S_\eta = \tau^{-1}\left(\{-\eta, \ldots, \eta\}^N\right) = \{y \in R \mid \|\tau(y)\|_\infty \leq \eta\}$, which corresponds to the set of polynomials in R with coefficients in $\{-\eta, \ldots, \eta\}$.

Another way of equipping K with a geometry is the canonical embedding. More precisely, for a number field K of degree N, let $\theta_1, \ldots, \theta_N$ denote the embeddings of K into $\mathbb{C}$. The canonical embedding $\theta \colon K \to \mathbb{C}^N$ is mapping $x \in K$ to $\theta(x) := (\theta_1(x), \ldots, \theta_N(x))^T$. It also helps us define the discriminant of the number field K as $\Delta_K = |\det(\theta_i(r_j))_{i,j}|^2$ for any basis (r_j) of R.

Let $\theta_i(\zeta_\nu) \in \mathbb{R}$ for $i \le t_1$ and for $t_1 < i \le t_2 + t_1$: $\theta_i(\zeta_\nu) \in \mathbb{C}$ and $\theta_{t_2+i}(\zeta_\nu) = \overline{\theta_i(\zeta_\nu)}$. Define

$$\mathbf{U}_H = \frac{1}{\sqrt{2}} \begin{pmatrix} \sqrt{2}\mathbf{I}_{t_1} & 0 & 0 \\ 0 & \mathbf{I}_{t_2} & \mathbf{I}_{t_2} \\ 0 & -i\mathbf{I}_{t_2} & i\mathbf{I}_{t_2} \end{pmatrix}.$$

Then $\theta_H(x) \coloneqq \mathbf{U}_H \cdot \theta(x) \in \mathbb{R}^N$. This embedding $\theta_H : K \to \mathbb{R}^N$ is sometimes called the Minkowski embedding. Since $\mathbf{U}_H$ is a unitary transformation, it does not affect the geometric lattice properties we consider in this paper. Hence, we will abuse the notation and call it canonical embedding for simplicity.

Note that the two geometries induced by the coefficient and canonical embeddings are in general not the same, but we can measure their distortion using the Vandermode matrix. Let $(\alpha_i)_{i\in\{1,\ldots,N\}} = (\theta_i(\zeta))_{i\in\{1,\ldots,N\}}$ be the roots of the defining polynomial $\Phi(X)$ of $K = \mathbb{Q}(\zeta)$. They form the Vandermonde matrix $\mathbf{V}_\Phi$ defined as

$$\mathbf{V}_\Phi = \begin{pmatrix} 1 & \alpha_1 & \cdots & \alpha_1^{N-1} \\ 1 & \alpha_2 & \cdots & \alpha_2^{N-1} \\ \vdots & \vdots & & \vdots \\ 1 & \alpha_N & \cdots & \alpha_N^{N-1} \end{pmatrix}.$$

Then $\forall x \in K : \theta(x) = \mathbf{U}_H \cdot \mathbf{V}_\Phi \cdot \tau(x)$. Denote as $\mathfrak{s}_1(\nu)$ the spectral norm of the Vandermonde matrix $\mathbf{V}_\Phi$ of the ν-th cyclotomic field, i.e., the largest singular value of $\mathbf{V}_\Phi$. Moreover, let $\mathfrak{s}_N(\nu)$ denote the smallest singular value of $\mathbf{V}_\Phi$, which is also the inverse of the spectral norm of $(\mathbf{V}_\Phi)^{-1}$. Since $\mathfrak{s}_1(\mathbf{U}_H) = \mathfrak{s}_N(\mathbf{U}_H) = 1$, for every $x \in K$ it holds

$$\|\tau(x)\| \cdot \mathfrak{s}_N(\nu) \le \|\theta(x)\| \le \|\tau(x)\| \cdot \mathfrak{s}_1(\nu). \tag{3}$$

In the case of power-of-two cyclotomics, i.e., $\nu = 2^{k+1}$ for a positive integer k and $N = 2^k$, we know that $\mathfrak{s}_1(\nu) = \mathfrak{s}_N(\nu) = \sqrt{N}$. Thus, τ and θ are isometric up to a factor of $\sqrt{N}$. The distortion between embeddings for other number fields was studied further in [BC22] and [RSW18]. For cyclotomic fields, we can lower and upper bound the spectral norm of $\mathbf{V}_\Phi$.

Lemma 1 ([ACX19, Proposition 2]). *For all positive integers ν*

$$\sqrt{\varphi(\nu)} \le \mathfrak{s}_1(\nu) \le \sqrt{t(\nu)},$$

where $t(\nu) = \nu$ when ν is odd and $t(\nu) = \nu/2$ when it is even. The inequality turns into an equality on both sides when ν is a power of 2.

The field discriminant can be bounded as follows.

Lemma 2 ([LPR13]). *For $\nu \in \mathbb{N}$, let $K = \mathbb{Q}(\zeta_\nu)$ be the ν-th cyclotomic number field. Denote $N = \varphi(\nu)$. Then*

$$|\Delta_K| = \frac{\nu^N}{\prod_{d|\nu} d^{\frac{N}{d-1}}} \le N^N.$$

Using the fact that every cyclotomic field K is a Galois field extension of $\mathbb{Q}$, we have a good grasp of the splitting behavior of the minimal polynomial $\Phi(X)$ into irreducible factor over finite fields. More precisely, for unramified prime q (this is guaranteed if $q \nmid \Delta_K$), the polynomial $\Phi(X)$ uniquely factors as $\Phi(X) = \prod_{j=1}^{f} \Phi_j(X) \bmod q$, such that $f|N$, all $\Phi_j(X)$ are distinct, have degree $\delta|N$ and are irreducible modulo q, where $N = f\delta$. Using that for cyclotomic fields, $R = \mathbb{Z}[X]/\langle \Phi(X)\rangle$, the above implies that for any q prime not dividing Δ_K, the quotient ring $R_q = R/qR$ of degree N factors into distinct finite fields as

$$R_q = \mathbb{Z}_q[X]/\langle \Phi(X)\rangle = \prod_{j=1}^{f} \mathbb{Z}_q[X]/\ \langle \Phi_j(X)\rangle.$$

We call R_q fully-splitting if $f = N$, high-splitting if δ is a small constant and low-splitting if f is a small constant.

If putting significantly stronger constraints on the relationship between q and $\Phi(X)$, the Φ_j's are guaranteed to have a special shape. This special shape in turn guarantees that short ring elements are invertible.

Lemma 3 ([LS18, Theorem 1.1]). *Let $K = \mathbb{Q}[X]/\langle \Phi(X)\rangle$ be the ν-th cyclotomic field of degree $N = \varphi(\nu)$, with R its ring of integers. Further, let $\nu = \prod p_i^{e_i}$ for $e_i \geq 1$ and let $\mu = \prod p_i^{f_i}$ for $1 \leq f_i \leq e_i$. If q is a prime such that $q = 1 \bmod \mu$ and the multiplicative order modulo ν of q is ν/μ, then the polynomial $\Phi(X)$ factors modulo q as $\Phi(X) = \prod_{j=1}^{\varphi(\mu)}(X^{\nu/\mu} - r_j) \bmod q$, for distinct $r_j \in \mathbb{Z}_q^{\times}$, where $X^{\nu/\mu} - r_j$ are irreducible in the ring $\mathbb{Z}_q[X]$. Further, for any element y of R_q satisfying $0 < \|\tau(y)\|_\infty < q^{1/\varphi(\mu)}/\mathfrak{s}_1(\mu)$, it holds that $y \bmod qR$ is a unit in R_q.*

An integral ideal $\mathcal{I}$ of a ring R is a subgroup of R such that $\mathcal{I} \cdot R = \mathcal{I}$. The definition also applies to the quotient ring $R_q = R/qR$ for an integer modulus q. We denote a principal ideal generated by $a \in R$ as $\langle a\rangle = aR$.

Lemma 4 (Ideals in R_q). *Consider a cyclotomic number field K, its ring of integers R and an unramified prime number q. Let $\Phi(X) = \prod_{j=1}^{f} \Phi_j(X) \bmod q$ be the factorisation of the corresponding cyclotomic polynomial in R_q. Then every ideal $\mathcal{I} \subsetneq R_q$ is of the form*

$$\mathcal{I} = \langle \prod_{j\in G} \Phi_j(X)\rangle, \text{ with } N(\mathcal{I}) = q^{\delta|G|},$$

for some subset $G \subset \{1, \ldots, f\}$. The set $G = \emptyset$ corresponds to the ideal $\langle 1\rangle$.

Proof. By the ideal correspondence theorem of commutative algebra, every ideal $\mathcal{I} \subset R_q$ corresponds to an ideal of $\mathbb{Z}_q[X]$ containing $\langle \Phi(X)\rangle$. Since $\mathbb{Z}_q$ is a field, $\mathbb{Z}_q[X]$ is a principal ideal domain. Then every ideal containing $\langle \Phi(X)\rangle$ is generated by a single polynomial $f(X)$ dividing $\Phi(X)$ modulo q. Hence, $\mathcal{I} = \langle \prod_{i\in G}(X^{\delta} - r_j)\rangle$ for some subset of indices $G \subset \{1, \ldots, f\}$. □

The norm of an integral ideal $\mathcal{I}$ of R is defined as $N(\mathcal{I}) = |R/\mathcal{I}|$. For any sub-ideal $\mathcal{I}' \subseteq \mathcal{I} \subseteq R$ the norm satisfies $N(\mathcal{I}) = |(R/\mathcal{I}')/(\mathcal{I}/\mathcal{I}')|$. In particular, for any $\mathcal{I}$ that contains qR we have $N(\mathcal{I}) = |R_q|/|\mathcal{I} \bmod qR|$.

The set of ideals $\mathcal{J} \subset R_q = R/qR$ are in one to one correspondence with ideals $\mathcal{I} \subsetneq R$ containing qR via a map $\mathcal{I} \bmod qR = \mathcal{J}$. Alternatively, we write $\mathcal{I} = qR + \mathcal{J}$ taking an arbitrary representative of $\mathcal{J}$ in R.

For any element a in the ring we define its algebraic norm as $N(a) \coloneqq \prod_{i=1}^{N} \theta_i(a)$, it verifies $|N(a)| = N(\langle a\rangle)$. Then for the fractional ideals $\frac{1}{a}\mathcal{I}$ where $a \in R$, $\mathcal{I} \subseteq R$ we define the norm as $N(\frac{1}{a}\mathcal{I}) \coloneqq \frac{N(\mathcal{I})}{|N(a)|}$. We denote R^* the dual of R as an R-module. The following lemma allows us to compute the norms of ideals inside R^*.

Lemma 5 ([LPR13, Section 2.5.4]). *Let K be a cyclotomic number field and R be its ring of integers. Then there exists an element $t \in R$ such that $R^* = \frac{1}{t} \cdot R$. Its norm is $R^* = |N(t)|^{-1} = \Delta_K^{-1}$.*

Then for and ideal $\mathcal{J} \subseteq R^*$ we have $N(\mathcal{J}) = N(t \cdot \mathcal{J}) \cdot \Delta_K^{-1}$ where $t \cdot \mathcal{J}$ is an integral ideal in R. We upper bound the sum of norms of all ideals in R_q.

Lemma 6 ([LPR13, Claim 7.3] adapted). *Let K be a cyclotomic number field of degree N and R its ring of integers, let $k > 1$ be an integer and $q \geq 2$ an unramified prime number. Let $R_q = R/qR$ factor into f fields of size q^δ, for some $f \cdot \delta = N$. Let $\mathbb{I}$ be the set of all integer ideals of R_q. Then*

$$\sum_{\mathcal{I} \in \mathbb{I}} |\mathcal{I}|^k = \sum_{\mathcal{I} \in \mathbb{I}} N(\mathcal{I})^k \leq \exp(1) \cdot q^{Nk}.$$

Additionally assuming $2f \cdot q^{-\delta k} \leq 1$, it holds

$$\sum_{\mathcal{I} \in \mathbb{I}} |\mathcal{I}|^{-k} = \sum_{\mathcal{I} \in \mathbb{I}} N(\mathcal{I})^k \leq 1 + 2f \cdot q^{-\delta k}.$$

Proof. Let $R_q = \prod_{j=1}^{f} \mathbb{Z}_q[X]/\langle \Phi_j(X) \rangle$ be the factorisation of the ring. Since q is an unramified prime by Lemma 4 every ideal of R_q is of the form $\mathcal{I} = \prod_{j \in G} \mathbb{Z}_q[X]/\langle \Phi_j(X) \rangle$ for some $G \subset \{1, \ldots, f\}$. Then for $\mathcal{I}' = \prod_{j \notin G} \mathbb{Z}_q[X]/\langle \Phi_j(X) \rangle$ it holds that $|\mathcal{I}| = N(\mathcal{I}')$. Hence, we can reorder the sum over all ideals and obtain $\sum_{\mathcal{I} \in \mathbb{I}} |\mathcal{I}|^k = \sum_{\mathcal{I}' \in \mathbb{I}} N(\mathcal{I}')^k$. This way derive bounds for the size of ideals from the bounds for their norm. We now extend the proof of [LPR13, Claim 7.3] to negative powers,

$$\sum_{\mathcal{I} \in \mathbb{I}} N(I)^{-k} \leq (1 + q^{-\delta k})^f \leq 1 + 2f \cdot q^{-\delta k},$$

using the bound on binomial formula $(1 + x)^n \leq 1 + 2nx$ for $2nx \leq 1$.

□

2.2 Probability Theory

The cardinality of a finite set S is denoted by $|S|$. For a probability distribution $\mathcal{P}$ over S, we denote by $x \leftarrow \mathcal{P}$ the process of sampling $x \in S$ according to $\mathcal{P}$. The probability of sampling x according to $\mathcal{P}$ is denoted by $\mathcal{P}(x)$. By $\mathcal{U}(S)$, we denote the uniform distribution over S, which is defined as $\Pr[x = y | x \leftarrow \mathcal{U}(S)] = 1/|S|$ for every $y \in S$. The support of a probability distribution $\mathcal{P}$ over S is defined as $\mathrm{Supp}(\mathcal{P}) = \{x \in S \mid \mathcal{P}(x) \neq 0\}$.

The collision probability of $\mathcal{P}$ is defined as $\mathrm{Coll}(\mathcal{P}) = \Pr[x = x' | x, x' \leftarrow \mathcal{P}]$. By $\mathbb{E}(\mathcal{P})$ we denote the expected value of $\mathcal{P}$. If $\mathcal{P}$ is a discrete probability distribution over S, the collision probability equals $\mathrm{Coll}(\mathcal{P}) = \sum_{x \in S} \mathcal{P}(x)^2$ and the expected value equals $\mathbb{E}(\mathcal{P}) = \sum_{x \in S} x \cdot \mathcal{P}(x)$. The min-entropy of $\mathcal{P}$ is defined as $\mathrm{H}_\infty(\mathcal{P}) = -\log_2 (\max_{x \in S} \mathcal{P}(x))$. It is true that $\mathrm{Coll}(\mathcal{P}) \leq \sum_{x \in S} \mathcal{P}(x) \cdot \max_{y \in S} \mathcal{P}(y) = \max_{y \in S} \mathcal{P}(y)$, which implies $\mathrm{Coll}(\mathcal{P}) \leq 2^{-\mathrm{H}_\infty(\mathcal{P})}$. For two probability distributions $\mathcal{P}$ (over S) and $\mathcal{Q}$ (over T) we define the *average conditional collision probability* as

$$\widetilde{\mathrm{Coll}}(\mathcal{P} \mid \mathcal{Q}) := \mathbb{E}_{y \leftarrow \mathcal{Q}} \Pr[x = x' \mid x, x' \leftarrow \mathcal{P}, y].$$

We remark that there are different flavours of how to define average conditional collision probability in the literature [FB14]. We opted for this one as it allows for equalities when working with the Rényi divergence (cf. Lemma 8 below). If both $\mathcal{P}$ and $\mathcal{Q}$ are discrete, this equals $\widetilde{\mathrm{Coll}}(\mathcal{P} \mid \mathcal{Q}) = \sum_{y \in T} \mathcal{Q}(y) \sum_{x \in S} \Pr[\mathcal{P} = x \mid y]^2$. If $\mathcal{Q}$ is independent of $\mathcal{P}$, we recover the ordinary collision probability, i.e., $\widetilde{\mathrm{Coll}}(\mathcal{P} \mid \mathcal{Q}) = \mathrm{Coll}(\mathcal{P})$. The average conditional collision probability is upper bounded by the average conditional min-entropy introduced in [DORS08, Sec. 2.4],

$$\widetilde{\mathrm{H}_\infty}(\mathcal{P} \mid \mathcal{Q}) := -\log_2 \left(\mathbb{E}_{y \leftarrow \mathcal{Q}} \max_{x \in S} \Pr[\mathcal{P} = x \mid y] \right).$$

More precisely, $\widetilde{\mathrm{Coll}}(\mathcal{P} \mid \mathcal{Q}) \leq \mathbb{E}_{y \leftarrow \mathcal{Q}} \max_{x \in S} \Pr[\mathcal{P} = x \mid y] = 2^{-\widetilde{\mathrm{H}_\infty}(\mathcal{P}|\mathcal{Q})}$.

Let $\mathcal{P}$ and $\mathcal{Q}$ be two discrete probability distributions over a discrete domain S. Their statistical distance is defined as $\mathrm{SD}(\mathcal{P}, \mathcal{Q}) = 1/2 \sum_{x \in S} |\mathcal{P}(x) - \mathcal{Q}(x)|$. We call $\mathcal{P}$ and $\mathcal{Q}$ *statistically close* in the security parameter λ if $\mathrm{SD}(\mathcal{P}, \mathcal{Q}) = \mathsf{negl}(\lambda)$. If the statistical distance is bounded by $\varepsilon > 0$ we write $\mathcal{P} \approx_\varepsilon \mathcal{Q}$. The statistical distance to uniform, even in the presence of leakage, can be upper bounded in terms of the (average conditional) collision probability of $\mathcal{P}$ under leakage $\mathcal{Q}$.

Lemma 7 ([DORS08, Lem. 2.4] implicit). *Let $\mathcal{P}$ be a discrete probability distribution over a finite set S and $\mathcal{Q}$ be a discrete probability distribution over a finite set T. Then,*

i) $\mathrm{SD}(\mathcal{P}, \mathcal{U}(S)) \leq \frac{1}{2}\sqrt{|S| \cdot \mathrm{Coll}(\mathcal{P}) - 1}$,
ii) $\mathrm{SD}((\mathcal{P}, \mathcal{Q}), (\mathcal{U}(S), \mathcal{Q})) \leq \frac{1}{2}\sqrt{|S| \cdot \widetilde{\mathrm{Coll}}(\mathcal{P} \mid \mathcal{Q}) - 1}$.

Let $\mathcal{P}$ and $\mathcal{Q}$ be two discrete probability distributions such that $\mathrm{Supp}(\mathcal{P}) \subset \mathrm{Supp}(\mathcal{Q})$. Their Rényi divergence (of order 2 and in exponential notation) is defined as $\mathrm{RD}(\mathcal{P}; \mathcal{Q}) = \sum_{x \in \mathrm{Supp}(\mathcal{P})} \frac{\mathcal{P}(x)^2}{\mathcal{Q}(x)}$. We also define the Rényi divergence of infinite order as $\mathrm{RD}_\infty(\mathcal{P}; \mathcal{Q}) = \max_{x \in \mathrm{Supp}(P)} \frac{\mathcal{P}(x)}{\mathcal{Q}(x)}$. For any two distributions $\mathcal{P}, \mathcal{Q}$ it holds $\mathrm{RD}(\mathcal{P}; \mathcal{Q}) \leq \mathrm{RD}_\infty(\mathcal{P}; \mathcal{Q})$. This and other properties of the Rényi divergence, such as data processing inequality, have been shown in [vEH14]. We call $\mathcal{P}$ and $\mathcal{Q}$ *Rényi close* if $\mathrm{RD}(\mathcal{P}; \mathcal{Q}) \in \mathcal{O}(1)$. The Rényi divergence to the uniform distribution can be expressed in terms of average conditional collision probability. To the best of our knowledge, this is the first time that the connection between the Rényi divergence and the collision probability in the presence of leakage is made formal.

Lemma 8. *Let $\mathcal{P}$ be a discrete probability distribution over a finite set S and $\mathcal{Q}$ be a discrete probability distribution over a finite set T. Then,*

i) $\mathrm{RD}(\mathcal{P}; \mathcal{U}(S)) = |S| \cdot \mathrm{Coll}(\mathcal{P})$.
ii) $\mathrm{RD}((\mathcal{P}, \mathcal{Q}), (\mathcal{U}(S), \mathcal{Q})) = |S| \cdot \widetilde{\mathrm{Coll}}(\mathcal{P} \mid \mathcal{Q})$.

Proof. We prove the equality in the presence of leakage. The non-leakage version follows by $\widetilde{\mathrm{Coll}}(\mathcal{P} \mid \mathcal{Q}) = \mathrm{Coll}(\mathcal{P})$ if $\mathcal{Q}$ is independent of $\mathcal{P}$, the multiplicativity of the Rényi divergence, as well as $\mathrm{RD}(\mathcal{Q}, \mathcal{Q}) = 1$.

First, we observe that $\mathrm{Supp}(\mathcal{P}) \subset \mathrm{Supp}(\mathcal{U}(S))$ and hence the Rényi divergence is well-defined. Let $\mathcal{U} := \mathcal{U}(S)$ and $\mathrm{RD} := \mathrm{RD}((\mathcal{P}, \mathcal{Q}); (\mathcal{U}, \mathcal{Q}))$. By the definition of the Rényi divergence and the average collision probability, it yields

$$\begin{aligned}
\mathrm{RD} :=& \sum_{(x,y) \in \mathrm{Supp}(\mathcal{P} \times \mathcal{Q})} \frac{\Pr\left[\mathcal{P} = x \wedge \mathcal{Q} = y\right]^2}{\Pr\left[\mathcal{U} = x \wedge \mathcal{Q} = y\right]} \\
&= |S| \cdot \sum_{(x,y) \in \mathrm{Supp}(\mathcal{P} \times \mathcal{Q})} \frac{\Pr\left[\mathcal{P} = x \mid \mathcal{Q} = y\right]^2 \Pr\left[\mathcal{Q} = y\right]^2}{\Pr\left[\mathcal{Q} = y\right]} \\
&= |S| \cdot \sum_{(x,y) \in \mathrm{Supp}(\mathcal{P} \times \mathcal{Q})} \Pr\left[\mathcal{P} = x \mid \mathcal{Q} = y\right]^2 \Pr\left[\mathcal{Q} = y\right] \\
&= |S| \cdot \sum_{(x,y) \in S \times T} \Pr\left[\mathcal{P} = x \mid \mathcal{Q} = y\right]^2 \Pr\left[\mathcal{Q} = y\right] \\
&= |S| \cdot \widetilde{\mathrm{Coll}}(\mathcal{P} \mid \mathcal{Q}),
\end{aligned}$$

where we used from line 1 to 2 that $\mathcal{U}$ and $\mathcal{Q}$ are independent. □

We recall a general result, lower bounding the average conditional collision probability by the min-entropy and the bit-size of the leakage.

Lemma 9 ([DORS08, Lem. 2.2]). *Let $\mathcal{P}$ and $\mathcal{Q}$ be two discrete probability distributions, where $\mathcal{Q}$ has at most $|T|$ possible values. Then, $\widetilde{\mathrm{H}_\infty}(\mathcal{P} \mid \mathcal{Q}) \geq \mathrm{H}_\infty(\mathcal{P}) - \log_2|T|$. This implies that $\widetilde{\mathrm{Coll}}(\mathcal{P} \mid \mathcal{Q}) \leq 2^{-\mathrm{H}_\infty(\mathcal{P})} \cdot |T|$.*

Let R be the ring of integers of a number field and q an unramified prime such that the factorisation of R_q into subfields is given by $R_q = \mathbb{Z}_q[X]/\langle \Phi(X)\rangle = \prod_{j=1}^{f} \mathbb{Z}_q[X]/\langle \Phi_j(X)\rangle$. For an element $x \in R_q$ and $j \in \{1, \ldots, f\}$, we call $x \bmod \Phi_j(X)$ the j-th CRT slot. An element x is invertible in R_q (denoted $x \in R_q^\times$) if and only if all of its CRT-slots are non-zero. We recall the notion of well-spread distributions over R_q as defined in [AAB$^+$24].

Definition 1 (Well-Spreadness). *Let $\mathcal{P}$ be a distribution over R_q and $B \in [0,1]$ be a constant. We say that $\mathcal{P}$ is B-well-spread if for all $j \in [f]$ and for all $y \in \mathbb{Z}_q[X]/\langle \Phi_j(X)\rangle$ it holds*

$$\Pr[x \bmod \Phi_j(X) = y | x \leftarrow \mathcal{P}] \leq B.$$

The closer B gets to $1/q^\delta$, the more $\mathcal{P}$ acts like the uniform distribution over each CRT-slot.

Lastly, we recall a useful lemma proven in [Mic07, Lem. 4.3] for ideals and generalised to modules in [BJRW23, Lem. B.3].

Lemma 10. *Let $\mathcal{A}$ be an arbitrary finite ring and n, m be positive integers. Consider an arbitrary vector $\mathbf{z} = (z_i)_{i=1}^m \in \mathcal{A}^m$. We define the ideal $\mathcal{I}_\mathbf{z} := \langle z_1, \ldots, z_m\rangle$. For $\mathbf{A} \leftarrow \mathcal{U}(\mathcal{A}^{n\times m})$, the vector $\mathbf{Az}$ is uniformly distributed over the module $\mathcal{I}_\mathbf{z}^n$. In particular, $\Pr_{\mathbf{A}\leftarrow\mathcal{U}(\mathcal{A}^{n\times m})}[\mathbf{Az} = \mathbf{0}] = \frac{1}{|\mathcal{I}_\mathbf{z}|^n}$.*

2.3 Lattices

A lattice Λ is a discrete additive subgroup of $\mathbb{R}^m$. It is usually defined as all integer linear combinations of a set of basis vectors that are linearly independent in $\mathbb{R}^m$. Formally, a set of $n \leq m$ basis vectors $\mathbf{b}_1, \ldots, \mathbf{b}_n \in (\mathbb{R}^m)^n$ defines a lattice $\Lambda(\mathbf{b}_1, \ldots, \mathbf{b}_n) := \{\sum_{i=1}^n \mathbf{b}_i x_i | x_i \in \mathbb{Z}\ \forall i\}$. Parameter m here is the lattice dimension and parameter n is its rank. In this work, we only work with full-rank lattices, that are lattices for which $n = m$ and hence dimension and rank can be used interchangeably. We further define the dual of a lattice Λ as $\Lambda^* = \{\mathbf{y} \in \mathrm{span}_\mathbb{R}(\Lambda) \colon \langle \mathbf{y}, \mathbf{x}\rangle \in \mathbb{Z}\ \forall \mathbf{x} \in \Lambda\}$. The first minimum of a lattice Λ with respect to the ℓ_2-norm is defined by $\lambda_1(\Lambda) = \min_{\mathbf{x}\in\Lambda\setminus\{\mathbf{0}\}} \|\mathbf{x}\|$. We write $\lambda_1^\infty(\Lambda)$ for the first minimum with respect to the infinity norm. Denote $\mathcal{C}_n$ a closed ball of dimension n and radius 1. Then for $i > 1$ the i-th minimum of the lattice is defined as $\lambda_i(\Lambda) = \min\{r > 0 : \dim \mathrm{span}(\Lambda \cap r\mathcal{C}_n) \geq i\}$

Let $\mathbf{\Sigma} \in \mathbb{R}^{m\times m}$ be a symmetric positive-definite matrix and $\mathbf{c} \in \mathbb{R}^m$ and Λ be an m-dimensional lattice. For any vector $\mathbf{x} \in \mathbb{R}^m$ we define the Gaussian mass function $\rho_{\sqrt{\mathbf{\Sigma}},\mathbf{c}}(\mathbf{x})$ as follows: $\rho_{\sqrt{\mathbf{\Sigma}},\mathbf{c}}(\mathbf{x}) := \exp\left(-\pi(\mathbf{x}-\mathbf{c})^T\mathbf{\Sigma}^{-1}(\mathbf{x}-\mathbf{c})\right)$. Then, the discrete Gaussian distribution $\mathcal{D}_{\Lambda,\sqrt{\mathbf{\Sigma}},\mathbf{c}}$ over Λ is given by

$$\forall \mathbf{x} \in \Lambda : \mathcal{D}_{\Lambda,\sqrt{\mathbf{\Sigma}},\mathbf{c}}(\mathbf{x}) = \frac{\rho_{\sqrt{\mathbf{\Sigma}},\mathbf{c}}(\mathbf{x})}{\sum_{\mathbf{z}\in\Lambda}\rho_{\sqrt{\mathbf{\Sigma}},\mathbf{c}}(\mathbf{z})}.$$

When $\mathbf{c} = \mathbf{0}$ or $\Lambda = \mathbb{Z}^m$ we omit them from the notation. When $\mathbf{\Sigma} = \sigma^2\mathbf{I}_m$, we say the discrete Gaussian distribution is spherical and simply write $\mathcal{D}_{\Lambda,\sigma,\mathbf{c}}$.

An important quantity related to discrete Gaussians over lattices is the smoothing parameter [MR07]. For some $\varepsilon > 0$ and a lattice Λ, we define $\eta_\varepsilon(\Lambda)$ as the smallest $s > 0$ such that $\rho_{1/s}(\Lambda^*) \leq 1 + \varepsilon$.

The notion of the smoothing parameter was adapted to elliptical Gaussians for example in [Pei10]. For that, we define a partial ordering on positive semi-definite matrices. For matrices $\mathbf{\Sigma}_1, \mathbf{\Sigma}_2 \in \mathbb{R}^{m \times m}$ we say $\mathbf{\Sigma}_1 \geq \mathbf{\Sigma}_2$ iff $\mathbf{\Sigma}_1 - \mathbf{\Sigma}_2$ is positive semi-definite. As a shorthand for any $s \geq 0$ we say $\mathbf{\Sigma} \geq s$ iff $\mathbf{\Sigma} \geq s \cdot \mathbf{I}_m$. Then for a symmetric positive semi-definite matrix $\mathbf{\Sigma} \in \mathbb{R}^{m \times m}$ the inequality $\mathbf{\Sigma} \geq \eta_\varepsilon^2(\Lambda) := s^2$ correctly implies $\rho_{\sqrt{\mathbf{\Sigma}^{-1}}}(\Lambda^*) = \sum_{\mathbf{x} \in \Lambda^*} \exp(-\pi \cdot \mathbf{x}^T \mathbf{\Sigma} \mathbf{x}) \leq \sum_{\mathbf{x} \in \Lambda^*} \exp(-\pi \cdot \mathbf{x}^T s^2 \mathbf{I}_m \mathbf{x}) = \rho_{1/s}(\Lambda^*) \leq 1 + \varepsilon$.

Lemma 11 ([Pei08]). *For any $\varepsilon > 0$ and n-dimensional lattice Λ it holds*

$$\eta_\varepsilon(\Lambda) \leq \frac{\sqrt{\ln(2n(1 + 1/\varepsilon))/\pi}}{\lambda_1^\infty(\Lambda^*)}$$

Lemma 12 ([MR07, Lemma 3.3] adapted). *For any $\varepsilon \in (0, 1/2)$ and a rank-d lattice Λ it holds $\eta_\varepsilon(\Lambda) \leq \lambda_d(\Lambda) \cdot \sqrt{\ln(2d(1 + 1/\varepsilon))/\pi}$.*

Lemma 13 ([LPR13, Lemma 2.6]). *For any n-dimensional lattice Λ*

$$\eta_{2^{-2n}}(\Lambda) \leq \sqrt{n}/\lambda_1(\Lambda^*).$$

Lemma 14 ([GMPW20, Lemma 2.3]). *Let $m \geq n > 0$ and $\mathbf{\Sigma} \geq 0$. For any n-dimensional lattice Λ, center $\mathbf{c} \in \mathbb{R}^n$ and an injective map $\mathbf{V} \in \mathbb{R}^{m \times n}$*

$$\mathbf{V} \cdot \mathcal{D}_{\Lambda, \sqrt{\mathbf{\Sigma}}, \mathbf{c}} = \mathcal{D}_{\mathbf{V} \cdot \Lambda, \sqrt{\mathbf{V} \cdot \mathbf{\Sigma} \cdot \mathbf{V}^T}, \mathbf{V} \cdot \mathbf{c}}$$

The following result is usually referenced to [GPV08, Corollary 2.8], but was already implicit in [MR07]. The version for non-spherical discrete Gaussians we state is for instance proven in [Jeu24, Lemma 1.25].

Lemma 15. *Let Λ, Λ' be m-dimensional lattices such that $\Lambda' \subseteq \Lambda$, let $\varepsilon \in (0, 1)$, $\mathbf{\Sigma} \in \mathbb{R}^{m \times m}$ be symmetric positive-definite and $\mathbf{c} \in \mathbb{R}^m$. Let $(\mathcal{D}_{\Lambda, \sqrt{\mathbf{\Sigma}}, \mathbf{c}} \bmod \Lambda')$ denote the distribution on Λ/Λ' obtained by sampling a vector from $\mathcal{D}_{\Lambda, \sqrt{\mathbf{\Sigma}}, \mathbf{c}}$ and reducing the sample modulo Λ'. Then for any $\sqrt{\mathbf{\Sigma}} \geq \eta_\varepsilon(\Lambda')$ and $\forall \mathbf{x} \in \Lambda/\Lambda'$:*

$$\Pr[\mathbf{x} \leftarrow \mathcal{D}_{\Lambda, \sqrt{\mathbf{\Sigma}}, \mathbf{c}} \bmod \Lambda'] \in \left[\frac{1 - \varepsilon}{1 + \varepsilon}, \frac{1 + \varepsilon}{1 - \varepsilon}\right] \frac{1}{|\Lambda/\Lambda'|}.$$

Corollary 1. *Let Λ, Λ' be m-dimensional lattices such that $\Lambda' \subseteq \Lambda$, let $\varepsilon \in (0, 1)$, $\mathbf{\Sigma} \in \mathbb{R}^{m \times m}$ be symmetric positive-definite and $\mathbf{c} \in \mathbb{R}^m$. Then for any $\sqrt{\mathbf{\Sigma}} \geq \eta_\varepsilon(\Lambda')$*

$$\mathrm{SD}((\mathcal{D}_{\Lambda, \sqrt{\mathbf{\Sigma}}, \mathbf{c}} \bmod \Lambda'), \mathcal{U}(\Lambda/\Lambda')) \leq \varepsilon/(1 - \varepsilon).$$

Corollary 2 ([BLL⁺15, Lemma 2.10] adapted). *Let Λ, Λ' be m-dimensional lattices such that $\Lambda' \subseteq \Lambda$, let $\varepsilon \in (0, 1)$, $\mathbf{\Sigma} \in \mathbb{R}^{m \times m}$ be symmetric positive-definite and $\mathbf{c} \in \mathbb{R}^m$. Then for any $\sqrt{\mathbf{\Sigma}} \geq \eta_\varepsilon(\Lambda')$*

$$\mathrm{RD}((\mathcal{D}_{\Lambda, \sqrt{\mathbf{\Sigma}}, \mathbf{c}} \bmod \Lambda'); \mathcal{U}(\Lambda/\Lambda')) \leq \mathrm{RD}_\infty((\mathcal{D}_{\Lambda, \sqrt{\mathbf{\Sigma}}, \mathbf{c}} \bmod \Lambda'); \mathcal{U}(\Lambda/\Lambda')) \leq \frac{1 + \varepsilon}{1 - \varepsilon}.$$

We remark that [Pre17,GJK24] propose a different formulation of the above corollary, making use of the notion of relative error. However, their bound reflects an asymptotic behaviour, whereas the one we use holds in the worst case.

Lemma 16 (**[Pei08, Corollary 5.3] adapted**). *Let Λ be an m-dimensional lattice and $\sigma > 0$. Then, for $\mathbf{x} \leftarrow \mathcal{D}_{\Lambda,\sigma}$*

$$\Pr[\|\mathbf{x}\|_\infty \geq \sigma \cdot r] \leq 2m \cdot e^{-\pi r^2}.$$

Any R-module $\mathcal{M} \subset K^r$ of rank r defines module lattices $\tau(\mathcal{M})$ and $\theta(\mathcal{M})$ in $\mathbb{R}^{Nr}$ through the coefficient embedding τ or the canonical embedding θ. We write $\mathcal{D}^\tau_{\mathcal{M},\sqrt{\mathbf{\Sigma}},\mathbf{c}}$ (respectively $\mathcal{D}^\theta_{\mathcal{M},\sqrt{\mathbf{\Sigma}},\mathbf{c}}$) to denote the discrete Gaussian distribution obtained by embedding the module $\mathcal{M}$ and the center $\mathbf{c} \in R^r$ through τ (resp. θ). Any ideal $\mathcal{I}$ of degree-N ring R defines a lattice in $\mathbb{R}^N$ through the coefficient embedding τ or the canonical embedding θ. Such lattices are called ideal lattices.

Lemma 17 (**[LPR13, Lemma 2.14]**). *Let K be a number field of degree N. For any fractional ideal $\mathcal{I}$ of K*

$$\sqrt{N} \cdot N(\mathcal{I})^{1/N} \leq \lambda_1(\theta(\mathcal{I})).$$

By using Equation 3, we obtain

$$\sqrt{N} \cdot N(\mathcal{I})^{1/N} / \mathfrak{s}_1(\nu) \leq \lambda_1(\tau(\mathcal{I})).$$

Lemma 18 (**[PR07, Lemma 6.5] adapted**). *Let $\varepsilon = 2^{-N}$. For any ideal $\mathcal{I}$ in degree N ring R we have*

$$\eta_\varepsilon(\theta(\mathcal{I})) \leq (N(\mathcal{I}) \cdot |\Delta_K|)^{1/N}.$$

Lemma 19 (**[LPR13, Claim 7.1]**). *Let $r, \varepsilon > 0$. For any n-dimensional lattice Λ*

$$\rho_{1/r}(\Lambda) \leq \max\left(1, \left(\frac{\eta_\varepsilon(\Lambda^*)}{r}\right)^n\right)(1+\varepsilon)$$

In particular for any fractional ideal $\mathcal{I} \subset K$ we have

$$\rho_{1/r}(\theta(\mathcal{I})) \leq \max\left(1, N(\mathcal{I})^{-1} r^{-N}\right)(1 + 2^{-2N})$$
$$\rho_{1/r}(\tau(\mathcal{I})) \leq \max\left(1, N(\mathcal{I})^{-1} \mathfrak{s}_1(\nu)^N \cdot r^{-N}\right)(1 + 2^{-2N})$$

To obtain the bound for the ideals we apply Lemma 13 and Lemma 17 to the general statement. In this work, we consider the special class of q-ary lattices. For $\mathbf{A} \in R_q^{n \times m}$ we define $\Lambda_q^\perp(\mathbf{A}) = \{\mathbf{y} \in R^m \colon \mathbf{A}\mathbf{y} = \mathbf{0} \bmod q\}$, as well as $\Lambda_q(\mathbf{A}^T) = \left\{\mathbf{y} \in R^m \colon \exists \mathbf{s} \in R^n \text{ s.t } \mathbf{A}^T\mathbf{s} = \mathbf{y} \bmod q\right\}$. Both define an R-module of rank m, which after being embedded via τ or θ, defines a corresponding module lattice. For $a \in R_q = \mathbb{Z}_q[X]/\langle \Phi(X)\rangle$ define $\mathrm{Rot}(a) := [\tau(a) \mid \cdots \mid \tau(X^{N-1} \cdot a)] \in \mathbb{Z}_q^{N \times N}$ and for $\mathbf{A} = (a_{ij})_{i,j} \in R_q^{n \times m}$ define a block matrix $\mathrm{Rot}(\mathbf{A}) := (\mathrm{Rot}(a_{ij}))_{i,j} \in \mathbb{Z}_q^{Nn \times Nm}$. We note that for the coefficient embedding $\tau(\Lambda_q^\perp(\mathbf{A})) = \Lambda_q^\perp(\mathrm{Rot}(\mathbf{A}))$.

Lemma 20 (**[RSW18, Lemma 5.1] adapted**). *Let K be a number field with R its ring of integers. Further, let $q \geq 2$ be a modulus defining R_q. Let $\mathbf{A} \in R_q^{n \times m}$ for integers $0 < n < m$. Then with respect to the canonical embedding θ*

$$\theta(\Lambda_q^\perp(\mathbf{A}))^* = \theta\left((R^*)^m + \left\{1/q \cdot \mathbf{A}^T \cdot \mathbf{s} \mid \mathbf{s} \in (R^*)^n\right\}\right),$$

where R^* is the dual of R seen as an ideal. With respect to the coefficient embedding τ using the transformation to integer lattices we get

$$\tau(\Lambda_q^\perp(\mathbf{A}))^* = \Lambda_q^\perp(\mathrm{Rot}(\mathbf{A}))^* = 1/q \cdot \Lambda_q(\mathrm{Rot}(\mathbf{A})^T).$$

Lemma 21 (**[BJRW23, Lemma 2.6]**). *Let K be a cyclotomic field with R its ring of integers of degree N. Further, let q be a prime modulus such that $R_q = \prod_{j=1}^f \mathbb{Z}_q[X]/\langle \Phi_j(X)\rangle$ for distinct, irreducible $\Phi_j(X)$ of degree δ and $N = f\delta$. Let $n \leq m$ be integers. Then*

$$\Pr_{\mathbf{A} \leftarrow \mathcal{U}(R_q^{n \times m})}(\mathbf{A} \cdot R_q^m = R_q^n) \geq \prod_{i=0}^{n-1}\left(1 - \frac{1}{q^{\delta(m-i)}}\right)^f.$$

3 LHL via Counting Ideals

Let K be a number field with R its ring of integers. Let m, n, q be positive integers and let $\mathcal{P}$ be a probability distribution over R_q^m with $R_q = R/qR$. We define the distribution $\mathcal{P}'$ as follows: Sample $\mathbf{A} \leftarrow \mathcal{U}(R_q^{n\times m})$ and $\mathbf{x} \leftarrow \mathcal{P}$, then output $(\mathbf{A}\mathbf{x} \bmod q, \mathbf{A}) \in R_q^n \times R_q^{n\times m}$. To ease readability, we sometimes assume the modulo q operation without writing it explicitly. Our overall goal is to show that, under clearly specified conditions, the distribution $\mathcal{P}'$ is (statistically or Rényi) close to the uniform distribution over $R_q^n \times R_q^{n\times m}$.

We start by bounding the statistical distance and the Rényi divergence using algebraic properties of ideals over the ring R_q in Section 3.1. In Section 3.2 we subsequently generalise the result to allow for leakage y on the hash input $\mathbf{x}$. Lastly, in Section 3.3 we derive different corollaries which replace the algebraic properties by more concrete terms already studied in the literature, e.g. collision probability or well-spreadness. Those corollaries both recover flavours of the leftover hash lemma over cyclotomic rings previously proposed in the literature, as well as provide new interpretations. Note that the statements of the main results Theorem 4+5 are purely over the ring and hence agnostic to the choice of the embedding. They actually apply to any finite ring, not necessarily cyclotomic. In contrast, the probabilistic terms such as collision probability and well-spreadness in Section 3.3 require specifying which embedding is used.

In Section 5 we give bounds with detailed proofs on well-spreadness and collision probability for various distributions over cyclotomic rings proposed in the lattice-based literature.

3.1 General Theorem

For $\mathbf{x} = (x_i)_{i=1}^m, \mathbf{x}' = (x_i')_{i=1}^m \in R_q^m$, we define the ideal $\mathcal{I}_{\mathbf{x},\mathbf{x}'}$ as $\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \langle x_1 - x_1', \ldots, x_m - x_m' \rangle \subset R_q$. This ideal has already implicitly appeared in [Mic07]. Sometimes, it is called the greatest common divisor ideal [Con]. Applying the one-to-one correspondence between ideals in R_q and ideals in R containing $\langle q \rangle$, one can see that the ideal $\mathcal{I}_{\mathbf{x},\mathbf{x}'}$ matches the notion of maximal belonging used in [LW20]. That is, $\mathbf{x} - \mathbf{x}'$ maximally belongs to the ideal $\mathcal{I} \subseteq R$ iff $\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I}'$, where $\mathcal{I}' \subseteq R_q$ is the unique ideal corresponding to $\mathcal{I}$.

Theorem 4. *Let K be a number field with R its ring of integers. Let m, n, q be positive integers and $\mathcal{P}$ be a distribution over R_q^m. We define*

$$\mathcal{P}' = \left\{(\mathbf{A}\mathbf{x} \bmod q, \mathbf{A}) \mid \mathbf{A} \leftarrow \mathcal{U}(R_q^{n\times m}), \mathbf{x} \leftarrow \mathcal{P}\right\}, \text{ and}$$
$$\mathcal{U}' = \left\{(\mathbf{u}, \mathbf{A}) \mid \mathbf{A} \leftarrow \mathcal{U}(R_q^{n\times m}), \mathbf{u} \leftarrow \mathcal{U}(R_q^n)\right\}.$$

Denote $\mathbb{I}$ the set of all ideals in R_q. Then, it holds

$$\mathrm{SD}(\mathcal{P}', \mathcal{U}') \le \frac{1}{2}\sqrt{\sum_{\substack{\mathcal{I}\in\mathbb{I}\\ \mathcal{I}\neq R_q}} N(\mathcal{I})^n \cdot \Pr_{\mathbf{x},\mathbf{x}'\leftarrow\mathcal{P}}[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I}]},$$

$$\mathrm{RD}(\mathcal{P}'; \mathcal{U}') = \sum_{\mathcal{I}\in\mathbb{I}} N(\mathcal{I})^n \cdot \Pr_{\mathbf{x},\mathbf{x}'\leftarrow\mathcal{P}}[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I}].$$

We highlight that the formula is exact in the Rényi divergence case.

Proof. Our goal is to bound $\mathrm{Coll}(\mathcal{P}')$, and then apply Lemma 7 and Lemma 8. Over the randomness of $\mathbf{A}, \mathbf{A}' \leftarrow \mathcal{U}(R_q^{n\times m})$ and $\mathbf{x}, \mathbf{x}' \leftarrow \mathcal{P}$ it yields

$$\begin{aligned}\mathrm{Coll}(\mathcal{P}') &= \Pr\left[\mathbf{A} = \mathbf{A}' \wedge \mathbf{A}\cdot\mathbf{x} = \mathbf{A}'\cdot\mathbf{x}'\right]\\ &= \Pr\left[\mathbf{A} = \mathbf{A}'\right] \cdot \Pr\left[\mathbf{A}\cdot\mathbf{x} = \mathbf{A}\cdot\mathbf{x}' | \mathbf{A} = \mathbf{A}'\right]\\ &= \frac{1}{|R_q|^{n\cdot m}} \cdot \Pr\left[\mathbf{A}(\mathbf{x}-\mathbf{x}') = \mathbf{0}\right].\end{aligned}$$

Let $\mathbf{x} = (x_1, \ldots, x_m)$ and $\mathbf{x}' = (x'_1, \ldots, x'_m)$. We condition the remaining term on the value of ideal $\mathcal{I}_{\mathbf{x},\mathbf{x}'}$ generated by $\mathbf{x}$ and $\mathbf{x}'$, thus

$$\begin{aligned}\Pr\big[\mathbf{A}(\mathbf{x}-\mathbf{x}') = \mathbf{0}\big] &= \sum_{\mathcal{I}\in\mathbb{I}} \Pr\big[\mathbf{A}(\mathbf{x}-\mathbf{x}') = \mathbf{0} | \mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I}\big] \cdot \Pr\big[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I}\big] \\ &= \sum_{\mathcal{I}\in\mathbb{I}} \frac{\Pr\big[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I}\big]}{|\mathcal{I}|^n}.\end{aligned}$$

The last equality holds by Lemma 10. To obtain the final statement for the Rényi divergence, we put everything together and use Lemma 8 with $\mathrm{Supp}(\mathcal{P}') \subseteq R_q^n \times R_q^{n\times m} := S$, where $|S| = |R_q|^{n\cdot(m+1)}$. It holds

$$\begin{aligned}\mathrm{RD}(\mathcal{P}';\mathcal{U}') &= |R_q|^{n\cdot(m+1)} \cdot \mathrm{Coll}(\mathcal{P}') \\ &= |R_q|^{n\cdot(m+1)} \cdot \frac{1}{|R_q|^{n\cdot m}} \cdot \sum_{\mathcal{I}\in\mathbb{I}} \frac{\Pr\big[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I}\big]}{|\mathcal{I}|^n} \\ &= \sum_{\mathcal{I}\in\mathbb{I}} N(\mathcal{I})^n \cdot \Pr\big[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I}\big].\end{aligned}$$

The last equality uses the fact that $N(\mathcal{I}) = |R_q|/|\mathcal{I}|$. The statement for the statistical distance is obtained analogously using Lemma 7, the fact that $N(R_q) = 1$ and that $\Pr\big[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = R_q\big] \leq 1$. □

3.2 General Theorem in the Presence of Leakage

In this section, besides the hash input distribution $\mathcal{P}$ over R_q^m, we also consider a leakage distribution $\mathcal{Q}$ over a domain T. The distribution $\mathcal{P}'$ is as before: Sample $\mathbf{A} \leftarrow \mathcal{U}(R_q^{n\times m})$ and $\mathbf{x} \leftarrow \mathcal{P}$, then output $(\mathbf{A}\mathbf{x} \bmod q, \mathbf{A})$. Our overall goal is to show that, under specific conditions, the distribution $\mathcal{P}'$ is (statistically or Rényi) close to the uniform distribution $\mathcal{U}(R_q^n \times R_q^{n\times m})$, even in the presence of leakage sampled from $\mathcal{Q}$, that depends on $\mathbf{x}$.

Recall the definition of the ideal $\mathcal{I}_{\mathbf{x},\mathbf{x}} = \langle x_1 - x'_1, \ldots, x_m - x'_m\rangle \subset R_q$ for $\mathbf{x} = (x_i)_{i=1}^m, \mathbf{x}' = (x'_i)_{i=1}^m \in R_q^m$. The statement of Theorem 5 below is similar to Theorem 4, where $\Pr_{\mathbf{x},\mathbf{x}'\leftarrow\mathcal{P}}[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I}]$ is replaced by $\mathbb{E}_{y\leftarrow\mathcal{Q}} \Pr_{\mathbf{x},\mathbf{x}'\leftarrow\mathcal{P}}[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I} \mid y]$.

Theorem 5. *Let K be a number field with R its ring of integers. Let m, n, q be positive integers, let $\mathcal{P}$ be a distribution over R_q^m and let $\mathcal{Q}$ be another distribution over a set T. We define*

$$\mathcal{P}' = \big\{(\mathbf{A}\mathbf{x} \bmod q, \mathbf{A}) \mid \mathbf{A} \leftarrow \mathcal{U}(R_q^{n\times m}), \mathbf{x} \leftarrow \mathcal{P}\big\}, \ \textit{and}$$

$$\mathcal{U}' = \big\{(\mathbf{u}, \mathbf{A}) \mid \mathbf{A} \leftarrow \mathcal{U}(R_q^{n\times m}), \mathbf{u} \leftarrow \mathcal{U}(R_q^n)\big\}.$$

Denote $\mathbb{I}$ the set of all ideals in R_q. Then, it holds

$$\mathrm{SD}\big((\mathcal{P}', \mathcal{Q}), (\mathcal{U}', \mathcal{Q})\big) \leq \frac{1}{2}\sqrt{\sum_{\substack{\mathcal{I}\in\mathbb{I} \\ \mathcal{I}\neq R_q}} N(\mathcal{I})^n \cdot \mathbb{E}_{y\leftarrow\mathcal{Q}} \Pr_{\mathbf{x},\mathbf{x}'\leftarrow\mathcal{P}}[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I} \mid y]},$$

$$\mathrm{RD}\big((\mathcal{P}', \mathcal{Q}); (\mathcal{U}', \mathcal{Q})\big) = \sum_{\mathcal{I}\in\mathbb{I}} N(\mathcal{I})^n \cdot \mathbb{E}_{y\leftarrow\mathcal{Q}} \Pr_{\mathbf{x},\mathbf{x}'\leftarrow\mathcal{P}}[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I} \mid y].$$

Proof. Similar to the proof of Theorem 4, we bound the average conditional collision probability of $\mathcal{P}'$ given leakage $\mathcal{Q}$ and then apply Lemma 7 and Lemma 8. Over the random choices of $\mathbf{A}, \mathbf{A}' \leftarrow \mathcal{U}(R_q^{n\times m})$ and $\mathbf{x}, \mathbf{x}' \leftarrow \mathcal{P}$ it holds

$$\begin{aligned}\widetilde{\mathrm{Coll}}(\mathcal{P}' \mid \mathcal{Q}) &= \mathbb{E}_{y\leftarrow\mathcal{Q}}\left[\Pr\left[(\mathbf{A}, \mathbf{A}\mathbf{x}) = (\mathbf{A}', \mathbf{A}'\mathbf{x}') \mid y\right]\right] \\ &= \mathbb{E}_{y\leftarrow\mathcal{Q}}\left[\Pr[\mathbf{A} = \mathbf{A}' \mid y] \cdot \Pr\left[\mathbf{A}\mathbf{x} = \mathbf{A}'\mathbf{x}' \mid y, \mathbf{A} = \mathbf{A}'\right]\right] \\ &= \mathbb{E}_{y\leftarrow\mathcal{Q}}\left[\Pr[\mathbf{A} = \mathbf{A}'] \cdot \Pr\left[\mathbf{A}\mathbf{x} = \mathbf{A}\mathbf{x}' \mid y\right]\right] \\ &= \frac{1}{|R_q|^{n\cdot m}} \cdot \mathbb{E}_{y\leftarrow\mathcal{Q}}\left[\Pr\left[\mathbf{A}(\mathbf{x} - \mathbf{x}') = \mathbf{0} \mid y\right]\right],\end{aligned}$$

where from line 1 to 2 we used the chain rule and from line 2 to 3 we used that $y \leftarrow \mathcal{Q}$ is independent of $\mathbf{A}, \mathbf{A}' \leftarrow \mathcal{U}(R_q^{n\times m})$.

Let $\mathbf{x} = (x_1, \ldots, x_m)$ and $\mathbf{x}' = (x_1', \ldots, x_m')$. We condition the term inside the expected value on the ideal $\mathcal{I}_{\mathbf{x},\mathbf{x}'}$ generated by the difference of $\mathbf{x}$ and $\mathbf{x}'$, that is

$$\begin{aligned}\Pr\left[\mathbf{A}(\mathbf{x} - \mathbf{x}') = \mathbf{0} \mid y\right] &= \sum_{\mathcal{I}\in\mathbb{I}} \Pr\left[\mathbf{A}(\mathbf{x} - \mathbf{x}') = \mathbf{0} \mid y \wedge \mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I}\right] \cdot \Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I} \mid y] \\ &= \sum_{\mathcal{I}\in\mathbb{I}} \frac{\Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I} \mid y]}{|\mathcal{I}|^n},\end{aligned}$$

where we used that Lemma 10 applies even in the presence of leakage y on $\mathbf{x}$ and $\mathbf{x}'$ as the result holds for arbitrary hash inputs. Using Lemma 7 and Lemma 8 together with $\mathrm{Supp}(\mathcal{P}') \subseteq R_q^{n\times m} \times R_q^n$ concludes the proof. For example, for the Rényi divergence, we have

$$\begin{aligned}\mathrm{RD}\left((\mathcal{P}', \mathcal{Q}); (\mathcal{U}', \mathcal{Q})\right) &= |R_q|^{n\cdot(m+1)} \cdot \widetilde{\mathrm{Coll}}(\mathcal{P} \mid \mathcal{Q}) \\ &= |R_q|^{n\cdot(m+1)} \cdot \frac{1}{|R_q|^{n\cdot m}} \cdot \mathbb{E}_{y\leftarrow\mathcal{Q}} \sum_{\mathcal{I}\in\mathbb{I}} \frac{\Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I} \mid y]}{|\mathcal{I}|^n} \\ &= \sum_{\mathcal{I}\in\mathbb{I}} \frac{|R_q|^n}{|\mathcal{I}|^n} \mathbb{E}_{y\leftarrow\mathcal{Q}} \Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I} \mid y] \\ &= \sum_{\mathcal{I}\in\mathbb{I}} N(\mathcal{I})^n \cdot \mathbb{E}_{y\leftarrow\mathcal{Q}} \Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I} \mid y].\end{aligned}$$

In the statistical distance we further use that $N(R_q) = 1$ and $\mathbb{E}_{y\leftarrow\mathcal{Q}} \Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = R_q \mid y] \leq 1$. □

3.3 Corollaries

The first corollary provides a LHL with respect to the (average conditional) collision probability. For a distribution $\mathcal{P}$ and ideal $\mathcal{I}$ of R_q, we denote by $\mathcal{P} \bmod \mathcal{I}$ the distribution which first samples an element from $\mathcal{P}$ and then computes its value modulo $\mathcal{I}$.

Corollary 3. *Let K be a number field with R its ring of integers. Let m, n, q be positive integers and $\mathcal{P}$ be a distribution over R_q^m, defining the distributions $\mathcal{P}'$ and $\mathcal{U}'$ as in Theorem 5. Let $\mathcal{Q}$ be the leakage distribution. Denote $\mathbb{I}$ the set of all ideals in R_q. Then, it holds*

$$\mathrm{SD}\left((\mathcal{P}', \mathcal{Q}), (\mathcal{U}', \mathcal{Q})\right) \leq \frac{1}{2}\sqrt{\sum_{\substack{\mathcal{I}\in\mathbb{I} \\ \mathcal{I}\neq R_q}} N(\mathcal{I})^n \cdot \widetilde{\mathrm{Coll}}(\mathcal{P} \bmod \mathcal{I} \mid \mathcal{Q})]},$$

$$\mathrm{RD}\left((\mathcal{P}', \mathcal{Q}); (\mathcal{U}', \mathcal{Q})\right) \leq \sum_{\mathcal{I}\in\mathbb{I}} N(\mathcal{I})^n \cdot \widetilde{\mathrm{Coll}}(\mathcal{P} \bmod \mathcal{I} \mid \mathcal{Q}).$$

If $\mathcal{Q}$ is independent of $\mathcal{P}$, it simplifies to

$$\mathrm{SD}(\mathcal{P}',\mathcal{U}') \leq \frac{1}{2}\sqrt{\sum_{\substack{\mathcal{I}\in\mathbb{I}\\ \mathcal{I}\neq R_q}} N(\mathcal{I})^n \cdot \mathrm{Coll}(\mathcal{P} \bmod \mathcal{I})},$$

$$\mathrm{RD}(\mathcal{P}';\mathcal{U}') \leq \sum_{\mathcal{I}\in\mathbb{I}} N(\mathcal{I})^n \cdot \mathrm{Coll}(\mathcal{P} \bmod \mathcal{I}).$$

Proof. Let $\mathbf{x} = (x_1, \ldots, x_m)$ and $\mathbf{x}' = (x_1', \ldots, x_m')$, defining $\mathcal{I}_{\mathbf{x},\mathbf{x}'}$. The corollary follows from Theorem 5 as for $\mathbf{x}, \mathbf{x} \leftarrow \mathcal{P}$ and $\mathcal{I} \in \mathbb{I}$, it holds

$$\begin{aligned}\mathbb{E}_{y\leftarrow\mathcal{Q}} \Pr\big[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I} \mid y\big] &\leq \mathbb{E}_{y\leftarrow\mathcal{Q}} \Pr[\forall i : x_i - x_i' \in \mathcal{I} \mid y] \\ &= \mathbb{E}_{y\leftarrow\mathcal{Q}}\mathrm{Coll}(\mathcal{P} \bmod \mathcal{I} \mid y) = \widetilde{\mathrm{Coll}}(\mathcal{P} \bmod \mathcal{I} \mid \mathcal{Q}).\end{aligned}$$

□

The corollary generalises the LHLs presented in [LW20, Thm. 5.7] and [LWZW24, Thm. 4] to arbitrary leakage on the hash input and recovers the recent LHL from [JLWG25, Lem. 5.1].

The LHL proven in [BI22, Lem. 12] can be seen as a special case of Corollary 3, where both N and q are powers of 2, the leakage takes at most $|R_q|$ values and the hash input distribution $\mathcal{P}$ is a spherical discrete Gaussian (over a non-trivial lattice). The main part of the proof consists of providing a global bound H for the min-entropy modulo any ideal and using Lemma 9. To do so, they make use of the special shape of $\mathbb{I}$ when both N and q are powers of 2.

In the case of a uniform bounded hash input distribution, one can compute the sum of norms and collision probabilities, as given in the following corollary. It covers the LHL versions proved in [Mic07, Thm. 4.2], [SSTX09, Thm. 3.2], and [BJRW23, Lem. 2.8].

Corollary 4 (Bounded Uniform). *Let K be a cyclotomic field with R its ring of integers. Further, let q be an unramified prime modulus such that $R_q = \prod_{j=1}^f \mathbb{Z}_q[X]/\langle \Phi_j(X)\rangle$ for distinct, irreducible $\Phi_j(X)$ of degree δ. Let m, n, η be positive integers such that $\eta < q/2$. Define $S_\eta = \{x \in R \colon \|\tau(x)\|_\infty \leq \eta\}$, and let $\mathcal{P}$ be the uniform distribution over $S_\eta^m \subset R_q^m$, defining the distributions $\mathcal{P}'$, and $\mathcal{U}'$ as in Theorem 4. Then, it holds*

$$\mathrm{SD}(\mathcal{P}',\mathcal{U}') \leq \frac{1}{2}\sqrt{\left(\left(\frac{q^n}{(2\eta+1)^m}\right)^\delta + 1\right)^f - 1},$$

$$\mathrm{RD}(\mathcal{P}';\mathcal{U}') \leq \left(\left(\frac{q^n}{(2\eta+1)^m}\right)^\delta + 1\right)^f.$$

Note that one can simplify (but loosen) the bounds using

$$\left(\left(\frac{q^n}{(2\eta+1)^m}\right)^\delta + 1\right)^f \leq \left(\frac{q^n}{(2\eta+1)^m} + 1\right)^N.$$

Proof. We use the statement of Corollary 3 as a starting point. Recall from Lemma 4 that every ideal $\mathcal{I}$ of R_q is of the form $\mathcal{I} = \langle \prod_{j\in G} \Phi_j(X)\rangle$ with norm $N(\mathcal{I}) = q^{\delta|G|}$, for some subset of indices $G \subset \{1, \ldots, f\}$. Using Lemma 28,

$$\begin{aligned}\sum_{\mathcal{I}\in\mathbb{I}} N(\mathcal{I})^n \cdot \mathrm{Coll}(\mathcal{P} \bmod \mathcal{I}) &= \sum_{G\subset\{1,\ldots,f\}} q^{n\delta|G|} \cdot \frac{1}{(2\eta+1)^{m\delta|G|}} \\ &= \sum_{j=0}^f \binom{f}{j} \left(\frac{q^{n\delta}}{(2\eta+1)^{m\delta}}\right)^j = \left(\left(\frac{q^n}{(2\eta+1)^m}\right)^\delta + 1\right)^f\end{aligned}$$

□

Similarly, we obtain a corollary for discrete Gaussian distributions in the coefficient embedding, implicit in [JLWG25].

Corollary 5 (Discrete Gaussian). *Let K be a cyclotomic field with R its ring of integers. Further, let q be an unramified prime modulus such that $R_q = \prod_{j=1}^{f} \mathbb{Z}_q[X]/\langle \Phi_j(X)\rangle$ for distinct, irreducible $\Phi_j(X)$ of degree δ. Let m, n be positive integers, σ be a positive real and $\mathbf{c} \in R^m$, and set $\mathcal{P} = \mathcal{D}^{\tau}_{R^m,\sigma,\mathbf{c}}$, defining the distributions $\mathcal{P}'$ and $\mathcal{U}'$ as in Theorem 4. If $\sigma \le \frac{\sqrt{\pi}}{2} \cdot \frac{q}{\sqrt{\ln(4mN)}}$ it holds*

$$\mathrm{SD}(\mathcal{P}',\mathcal{U}') \le \frac{1}{\sqrt{2}} \sqrt{\left(\left(\frac{q^n}{\sigma^m}\right)^{\delta} + 1\right)^{f} - 1},$$

$$\mathrm{RD}(\mathcal{P}';\mathcal{U}') \le 2\left(\left(\frac{q^n}{\sigma^m}\right)^{\delta} + 1\right)^{f}.$$

Proof. The proof is very similar to the one of Corollary 4. Using Lemma 30 and that $\sigma \le \frac{\sqrt{\pi}}{2} \cdot \frac{q}{\sqrt{\ln(4mN)}} \le \frac{\sqrt{\pi}}{2} \cdot \frac{q}{\sqrt{\ln(4m\delta|G|)}}$ for all $G \subseteq \{1,\dots,f\}$, as well as $\mathrm{Coll}(\mathcal{P} \bmod \mathcal{I}) \le 2^{-\mathrm{H}_\infty(\mathcal{P} \bmod \mathcal{I})}$ it holds

$$\begin{aligned}
\sum_{\mathcal{I}\in\mathbb{I}} N(\mathcal{I})^n \cdot \mathrm{Coll}(\mathcal{P} \bmod \mathcal{I}) &\le \sum_{\mathcal{I}\in\mathbb{I}} N(\mathcal{I})^n \cdot 2^{-\mathrm{H}_\infty(\mathcal{P} \bmod \mathcal{I})} \\
&\le \sum_{G\subset\{1,\dots,f\}} q^{n\delta|G|} \cdot 2^{-(m\delta|G|\log\sigma - 1)} \\
&= 2\sum_{j=0}^{f} \binom{f}{j} \left(\frac{q^{n\delta}}{2^{m\delta\log\sigma}}\right)^{j} = 2\left(\left(\frac{q^n}{\sigma^m}\right)^{\delta} + 1\right)^{f},
\end{aligned}$$

concluding the proof for the Rényi divergence. For the statistical distance, we have a similar reasoning

$$\begin{aligned}
\sum_{\mathcal{I}\in\mathbb{I},\mathcal{I}\neq R_q} N(\mathcal{I})^n \cdot 2^{-\mathrm{H}_\infty(\mathcal{P} \bmod \mathcal{I})} &\le \sum_{G\subset\{1,\dots,f\},G\neq\emptyset} q^{n\delta|G|} \cdot 2^{-(m\delta|G|\log\sigma - 1)} \\
&= 2\left(\sum_{j=0}^{f} \binom{f}{j} \left(\frac{q^{n\delta}}{2^{m\delta\log\sigma}}\right)^{j} - 1\right) = 2\left(\left(\left(\frac{q^n}{\sigma^m}\right)^{\delta} + 1\right)^{f} - 1\right).
\end{aligned}$$

□

For general distributions, it might be difficult to compute the product of norms and collision probabilities for each ideal. Then, one can upper bound either the norms or the collision probabilities separately, as done in the two corollaries below.

The following statement generalises the LHL in [KY16, Lem. 4] from power-of-two cyclotomics to general cyclotomics, and additionally allows for arbitrarily shaped leakage. Essentially, it says that if R_q splits into few subfields and $\mathcal{P}$ provides elements of short norm with high probability, the ideal $\mathcal{I}_{\mathbf{x},\mathbf{x}'}$ generates the full ring unless $\mathbf{x}$ and $\mathbf{x}'$ collide. In other words, in the low-splitting regime, the function $\mathbf{A}\colon \mathbf{x} \mapsto \mathbf{A}\mathbf{x} \bmod q$ defines a universal hash function and we recover the standard (generalised) LHL [ILL89,HILL99,DORS08]. To the best of our knowledge, we are the first to derive a generalisation of the LHL in the low-splitting regime that allows for leakage.

Corollary 6 (Low-Splitting). *Let $K = \mathbb{Q}[X]/\langle \Phi(X) \rangle$ be the ν-th cyclotomic field of degree $N = \varphi(\nu)$, with R its ring of integers. Further, let $\nu = \prod p_i^{e_i}$ for $e_i \geq 1$ and let $\mu = \prod p_i^{f_i}$ for $1 \leq f_i \leq e_i$. Let q be a prime such that $q = 1 \bmod \mu$ and the multiplicative order modulo ν of q is ν/μ. Let $\mathfrak{s}_1(\mu)$ be the spectral norm of the Vandermonde matrix of $\mathbb{Q}(\zeta_\mu)$ and let η be a positive integer such that $2\eta < q^{1/\varphi(\mu)}/\mathfrak{s}_1(\mu)$, defining $S_\eta = \{x \in R_q \colon \|\tau(x)\|_\infty \leq \eta\}$. Further, let $\mathcal{P}$ be a distribution over R_q^m such that $\Pr[\forall i \in [m] : x_i \in S_\eta \mid \sum_{i\in[m]} x_i \cdot X^i = x \leftarrow \mathcal{P}] \geq 1 - C$ for some real $C \in [0,1]$. Moreover, let $\mathcal{Q}$ be the leakage distribution. We define the distributions $\mathcal{P}'$ and $\mathcal{U}'$ as in Theorem 5. Then, it holds*

$$\mathrm{SD}((\mathcal{P}',\mathcal{Q}),(\mathcal{U}',\mathcal{Q})) \leq \frac{1}{2}\sqrt{|R_q|^n \cdot \left(\widetilde{\mathrm{Coll}}(\mathcal{P} \mid \mathcal{Q}) + 2C\right)}$$

$$\mathrm{RD}((\mathcal{P}',\mathcal{Q});(\mathcal{U}',\mathcal{Q})) \leq 1 + |R_q|^n \cdot \left(\widetilde{\mathrm{Coll}}(\mathcal{P} \mid \mathcal{Q}) + 2C\right).$$

If $\mathcal{Q}$ is independent of $\mathcal{P}$, this simplifies to

$$\mathrm{SD}(\mathcal{P}',\mathcal{U}') \leq \frac{1}{2}\sqrt{|R_q|^n \cdot (\mathrm{Coll}(\mathcal{P}) + 2C)},$$

$$\mathrm{RD}(\mathcal{P}';\mathcal{U}') \leq 1 + |R_q|^n \cdot (\mathrm{Coll}(\mathcal{P}) + 2C).$$

Proof. First, we recall that $N(\mathcal{I}) \leq |R_q|$ for every ideal $\mathcal{I}$ of R_q. By the linearity of expectation, it is thus enough to show that for $\mathbf{x}, \mathbf{x}' \leftarrow \mathcal{P}$ it yields

$$\begin{aligned}\sum_{\substack{\mathcal{I}\in\mathbb{I}\\ \mathcal{I}\neq R_q}} \mathbb{E}_{y\leftarrow\mathcal{Q}} \Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I} \mid y] &= \mathbb{E}_{y\leftarrow\mathcal{Q}} \sum_{\substack{\mathcal{I}\in\mathbb{I}\\ \mathcal{I}\neq R_q}} \Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} = \mathcal{I} \mid y] \\ &= \mathbb{E}_{y\leftarrow\mathcal{Q}} \Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} \neq R_q \mid y] \leq \widetilde{\mathrm{Coll}}(\mathcal{P} \mid \mathcal{Q}) + 2C.\end{aligned}$$

By the law of total probabilities and the linearity of the expectation we have

$$\begin{aligned}\mathbb{E}_y \Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} \neq R_q \mid y] &= \mathbb{E}_y \big(\Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} \neq R_q \mid \mathbf{x} \neq \mathbf{x}', y] \cdot \Pr[\mathbf{x} \neq \mathbf{x}' \mid y] \\ &\quad + \Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} \neq R_q \mid \mathbf{x} = \mathbf{x}', y] \cdot \Pr[\mathbf{x} = \mathbf{x}' \mid y]\big) \\ &\leq \mathbb{E}_y(\Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} \neq R_q \mid \mathbf{x} \neq \mathbf{x}', y]) + \mathbb{E}_y(\mathrm{Coll}(\mathcal{P} \mid y)) \\ &= \mathbb{E}_y(\Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} \neq R_q \mid \mathbf{x} \neq \mathbf{x}', y]) + \widetilde{\mathrm{Coll}}(\mathcal{P} \mid \mathcal{Q}).\end{aligned}$$

It remains to bound the first probability term by $2C$. It holds

$$\begin{aligned}\Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} \neq R_q \mid \mathbf{x} \neq \mathbf{x}', y] &\geq \Pr[\mathcal{I}_{\mathbf{x},\mathbf{x}'} \neq R_q \mid \mathbf{x} \neq \mathbf{x}' \wedge \left\|\mathbf{x} - \mathbf{x}'\right\|_\infty \leq 2\eta, y] \\ &\quad + \Pr[\left\|\mathbf{x} - \mathbf{x}'\right\|_\infty > 2\eta] \geq 0 + 2C,\end{aligned}$$

where we used the properties of $\mathcal{P}$ and Lemma 3. Note that the latter result holds for arbitrary elements fulfilling the conditions, so even in the presence of leakage. This concludes the proof. □

Remark 1. To apply the corollary above in the presence of leakage one has different possibilities. Most available literature uses the chain rule from Lemma 9. Informally, it lower bounds the incurred entropy loss by the size of the leakage space. Alternatively, one can use more specific chain rules for concrete choices of distributions. As an example, [BD20, Sec. 5] proved a lower bound where, given $\mathbf{x} \leftarrow \mathcal{P}$ for arbitrary distribution $\mathcal{P}$, the distribution $\mathcal{Q}$ samples a Gaussian error $\mathbf{e}$ and outputs $\mathbf{x} + \mathbf{e}$.

Remark 2. Note that the above result enforces a relationship between the norm bound η of the hash input and the splitting behaviour of the ring R modulo q. In particular, the larger is μ (and thus $\mathfrak{s}_1(\mu)$ and $\varphi(\mu)$), the larger q needs to be for a fixed η to fulfil the condition. For example, if we are in the high-splitting regime, where μ/ν is a small constant, the modulus q would need to be exponentially large, even for the smallest $\eta = 1$. Thus, for practical parameter sets, this approach requires a low-splitting regime.

Next, we provide an alternative corollary, involving the well-spreadness of the hash input distribution $\mathcal{P}$. Even though it generally leads to looser bounds than Corollary 6, it can be used in the high-splitting regime with practical parameters. We highlight that to the best of our knowledge, we are the first to use well-spreadness in a context outside of compact lattice-based proof systems. In particular, we make a novel connection between well-spread distributions and randomness extraction. Note that we do not propose a leakage variant for it, as well-spreadness does depend on the distribution. Such generalisation would require a notion of well-spreadness under leakage, which we leave as future work.

Corollary 7 (Well-Spread). *Let K be a cyclotomic field such that R is its ring of integers. Let m, n, q be positive integers, where q is an unramified prime. For $B \in [0, 1]$, let $\mathcal{X}$ be a B-well-spread distribution over R_q, defining the distribution $\mathcal{P} = (\mathcal{X})^m$ and $\mathcal{P}'$ and $\mathcal{U}'$ as in Theorem 4. Then, it holds*

$$\mathrm{SD}(\mathcal{P}',\mathcal{U}') \leq \frac{1}{2}\sqrt{B^m \cdot \exp(1) \cdot |R_q|^n}, \quad \textit{and} \quad \mathrm{RD}(\mathcal{P}';\mathcal{U}') \leq B^m \cdot \exp(1) \cdot |R_q|^n.$$

Proof. Let $R_q = \prod_{j=1}^{f} \mathbb{Z}_q[X]/\langle \Phi_j(X)\rangle$ be the unique factorisation into subfields, and for $j \in [f]$, denote the j-th CRT slot of R_q as $\mathcal{I}_j = \mathbb{Z}_q[X]/\langle \Phi_j(X)\rangle$. By Lemma 4, ideals of R_q are of the form $\mathcal{I} = \prod_{j \in G} \mathcal{I}_j$ for $G \subset \{1, \ldots, f\}$. Then for arbitrary $j \in G$

$$\mathrm{Coll}(\mathcal{P} \bmod \mathcal{I}) = \Pr_{\mathbf{x},\mathbf{y} \leftarrow \mathcal{P}}[\mathbf{x} = \mathbf{y} \bmod \mathcal{I}] \leq \Pr_{\mathbf{x},\mathbf{y} \leftarrow \mathcal{P}}[\mathbf{x} = \mathbf{y} \bmod \mathcal{I}_j] \leq B^m.$$

Here the last inequality follows from the well-spreadness of $\mathcal{X}$ and $\mathcal{P}$ being a product of m independent distributions. By Lemma 6 for q prime, $\sum_{\substack{I \in \mathcal{I} \\ I \neq R_q}} N(\mathcal{I})^n \leq \sum_{I \in \mathcal{I}} N(\mathcal{I})^n \leq |R_q|^n \cdot \exp(1)$ and the corollary follows. □

4 LHL via Smoothing

In this section, we consider an alternative approach to obtain an LHL over cyclotomic rings. It focuses on the distribution of a discrete Gaussian vector $\mathbf{x}$ over R multiplied with a matrix $\mathbf{A}$ sampled from some distribution $\mathcal{P}_A$ over R_q. We prove that the output distribution is independent of $\mathcal{P}_A$ and close to uniform with respect to statistical distance and Rényi divergence. Note that in Section 3 we only looked at $\mathcal{P}_A$ being the uniform distribution. In this section, we take a more general perspective as slightly different distributions have been studied in the literature. The only works we are aware of that use the Rényi divergence in combination with a smoothing-based LHL (in slightly different contexts), are [BLR+18, Cor. 4.4] and [GJK24, Lem. 14].

We start by providing the general result in Section 4.1. It guarantees closeness to uniform as long as the width of $\mathbf{x}$ lies above the smoothing parameter of the kernel lattice of $\mathbf{A}$ and the image of $\mathbf{A}$ is of full rank, both with overwhelming probability over $\mathbf{A} \leftarrow \mathcal{P}_A$. Moreover, we show in Corollary 8 that once we can guarantee both conditions with respect to one embedding, we obtain an LHL for the other embedding too, hence allowing to switch between canonical and coefficient embedding easily. To the best of our knowledge, this observation is novel and leads to the first smoothing-based LHL in the coefficient embedding for arbitrary cyclotomic rings.

In Section 4.2, we give a concrete analysis of the smoothing parameter of the kernel lattice $\Lambda_q^\perp(\mathbf{A})$ for different choices of $\mathcal{P}_A$. The section both recalls existing results on non-uniform $\mathcal{P}_A$ and shows that we can generalise them to statements for $\mathbf{A}$ that comes from the uniform distribution. In Section 4.3 and Section 4.4 we look at smoothing in the presence of exact and noisy linear leakage, respectively.

4.1 General Theorem

In this section, we formally describe the general *smoothing-based* argument, leaving the analysis of the smoothing parameter for later. We emphasise that the proof itself is not novel, but separating the general argument from the concrete smoothing analysis allows for a more modular perspective. In particular, it helped us realise that we can switch between the embeddings (Corollary 8) and enabled a modular analysis of noisy linear leakage (Theorem 8).

Theorem 6. *Let K be a cyclotomic number field with R its ring of integers. For positive integer q, we define the finite ring R_q. Let $\mathcal{P}_A$ be a distribution on the set of matrices $R_q^{n\times m}$ for $n \le m$. Let $\varepsilon \in (0,1)$, and $\beta, \delta_1, \delta_2 > 0$ and $\pi \in \{\tau, \theta\}$ represent an arbitrary lattice embedding. Suppose that*

$$\Pr\Big[\eta_\varepsilon(\pi(\Lambda_q^\perp(\mathbf{A}))) > \beta | \mathbf{A} \leftarrow \mathcal{P}_A\Big] < \delta_1,$$
$$\Pr\big[\mathrm{im}(\mathbf{A}) \neq R_q^n | \mathbf{A} \leftarrow \mathcal{P}_A\big] < \delta_2.$$

Let $\mathbf{c} \in R^m$, $\mathbf{\Sigma} \in \mathbb{R}^{Nm\times Nm}$ positive definite s.t. $\mathbf{\Sigma} \ge \beta^2$. Define

$$\mathcal{P}' = \Big\{(\mathbf{A}\mathbf{x} \bmod q, \mathbf{A}) \mid \mathbf{A} \leftarrow \mathcal{P}_A, \mathbf{x} \leftarrow \mathcal{D}^\pi_{R^m,\sqrt{\mathbf{\Sigma}},\mathbf{c}}\Big\},$$
$$\mathcal{U}' = \big\{(\mathbf{u}, \mathbf{A}) \mid \mathbf{A} \leftarrow \mathcal{P}_A, \mathbf{u} \leftarrow \mathcal{U}(R_q^n)\big\}.$$

Then

$$\mathrm{SD}(\mathcal{P}',\mathcal{U}') < \varepsilon/(1-\varepsilon) + \delta_1 + \delta_2, \qquad \mathrm{RD}(\mathcal{P}';\mathcal{U}') < \frac{1+\varepsilon}{1-\varepsilon} + (\delta_1+\delta_2)\cdot q^{Nn}.$$

Proof. We partition the set $R_q^{n\times m}$ into $\mathcal{S} = \{\mathbf{A} \in R_q^{n\times m} | \eta_\varepsilon(\pi(\Lambda_q^\perp(\mathbf{A}))) \le \beta, \mathrm{im}(\mathbf{A}) = R_q^n\}$, that is, matrices that satisfy both constrains, and its complement. We first prove the statements for an arbitrary fixed matrix $\mathbf{A} \in \mathcal{S}$.

Algebraically, the mapping $\mathbf{A} : \mathbf{x} \in R^m \mapsto \mathbf{A}\mathbf{x} \bmod q \in R_q^n$ is an isomorphism between $R^m/\ker(\mathbf{A})$ and $\mathrm{im}(\mathbf{A})$. Note that $\ker(\mathbf{A}) = \Lambda_q^\perp(\mathbf{A})$ and $\mathrm{im}(\mathbf{A}) = R_q^n$. Therefore, the image $\mathbf{A}\mathbf{x} \bmod q$ is uniform in R_q^n if and only if $\mathbf{x} \bmod \Lambda_q^\perp(\mathbf{A})$ is uniform in $R^m/\Lambda_q^\perp(\mathbf{A})$.

Since $\mathbf{A} \in \mathcal{S}$, we know that $\mathbf{x} \leftarrow \mathcal{D}^\pi_{R^m,\sqrt{\mathbf{\Sigma}},\mathbf{c}}$ has the standard deviation $\mathbf{\Sigma} \ge \beta^2$ exceeding the smoothing parameter of $\pi(\Lambda_q^\perp(\mathbf{A}))$. Then Corollary 1 for lattices $\Lambda = \pi(R^m)$ and $\Lambda' = \pi(\Lambda_q^\perp(\mathbf{A}))$ implies

$$\mathrm{SD}(\mathbf{A}\mathbf{x} \bmod q, \mathcal{U}(R_q^n)) = \mathrm{SD}(\mathbf{x} \bmod \Lambda_q^\perp(\mathbf{A}), \mathcal{U}(R^m/\Lambda_q^\perp(\mathbf{A}))) \le \varepsilon/(1-\varepsilon).$$

Similarly, by Corollary 2 for the Rényi divergence

$$\mathrm{RD}(\mathbf{A}\mathbf{x} \bmod q; \mathcal{U}(R_q^n)) = \mathrm{RD}(\mathbf{x} \bmod \Lambda_q^\perp(\mathbf{A}); \mathcal{U}(R^m/\Lambda_q^\perp(\mathbf{A}))) \le \frac{1+\varepsilon}{1-\varepsilon}.$$

Now consider a matrix $\mathbf{A}$ sampled at random. The probability that $\mathbf{A} \leftarrow \mathcal{P}_A$ is not in $\mathcal{S}$ is at most $\delta_1 + \delta_2$. The statistical distance changes as

$$\mathrm{SD}(\mathcal{P}',\mathcal{U}') = \mathbb{E}_{\mathbf{A}}(\mathrm{SD}(\mathbf{A}\mathbf{x} \bmod q, \mathcal{U}(R_q^n))) \le \varepsilon/(1-\varepsilon) + \delta_1 + \delta_2,$$

because $\mathrm{SD}(\mathbf{Ax}, \mathcal{U}(R_q^n)) \leq 1$. By definition, for the Reńyi divergence we have

$$\begin{aligned}\mathrm{RD}(\mathbf{Ax}; \mathcal{U}(\operatorname{im}\mathbf{A})) &\leq |\operatorname{im}\mathbf{A}| \cdot \sum_{\mathbf{v} \in \operatorname{im}\mathbf{A}} \Pr\Big[\mathbf{v} = \mathbf{Ax} | \mathbf{x} \leftarrow \mathcal{D}^{\pi}_{R^m, \sqrt{\mathbf{\Sigma}}, \mathbf{c}}\Big]^2 \\ &\leq |\operatorname{im}\mathbf{A}| \leq |R_q^n|.\end{aligned}$$

Hence $\mathrm{RD}(\mathcal{P}'; \mathcal{U}') =$

$$\begin{aligned}&= \sum_{\mathbf{A} \in \mathcal{S}} \Pr[\mathbf{A} \leftarrow \mathcal{P}_A] \cdot \mathrm{RD}(\mathbf{Ax}; \mathcal{U}(\operatorname{im}\mathbf{A})) + \sum_{\mathbf{A} \notin \mathcal{S}} \Pr[\mathbf{A} \leftarrow \mathcal{P}_A] \cdot \mathrm{RD}(\mathbf{Ax}; \mathcal{U}(\operatorname{im}\mathbf{A})) \\ &\leq \frac{1+\varepsilon}{1-\varepsilon} + \Pr[\mathbf{A} \notin \mathcal{S} | \mathbf{A} \leftarrow \mathcal{P}_A] \cdot |R_q^n| = \frac{1+\varepsilon}{1-\varepsilon} + (\delta_1 + \delta_2) \cdot q^{Nn}.\end{aligned}$$

□

The theorem statement holds for both the canonical embedding θ and the coefficient embedding τ of the module lattices, as long as we have good bounds for the smoothing parameter in the respective kernel lattice. We demonstrated it for cyclotomic fields, but the proof actually generalises to any number field, as long as we know how to do the smoothing analysis.

We show that it is enough to know a bound for a smoothing parameter in one embedding, as it already allows us to apply the theorem to the other embedding using the Vandermonde transformation map between embeddings.

Corollary 8. *For $\nu \geq 1$ let R be the ν-th cyclotomic ring of degree N. Suppose that the conditions of Theorem 6 on distribution $\mathcal{P}_A$ are satisfied in the canonical embedding θ, then the statistical and Rényi closeness holds for $\mathbf{x} \leftarrow \mathcal{D}^{\tau}_{R^m, \sqrt{\mathbf{\Sigma}}, \mathbf{c}}$ with $\mathbf{\Sigma} \geq \beta^2 / \mathfrak{s}_N(\nu)^2$. Similarly, if the conditions are satisfied in the coefficient embedding τ, they imply the statement for $\mathbf{x} \leftarrow \mathcal{D}^{\theta}_{R^m, \sqrt{\mathbf{\Sigma}}, \mathbf{c}}$ with $\mathbf{\Sigma} \geq \beta^2 \cdot \mathfrak{s}_1(\nu)^2$.*

Proof. Let $\mathbf{V} = \mathbf{I}_m \otimes (\mathbf{U}_H \cdot \mathbf{V}_{\Phi}) \in \mathbb{R}^{Nm \times Nm}$ be the embedding distortion map. Then for any subset $\Lambda \subset R^m$ we have $\mathbf{V} \cdot \tau(\Lambda) = \theta(\Lambda)$ and $\mathbf{V}^{-1} \cdot \theta(\Lambda) = \tau(\Lambda)$. Since $\mathbf{U}_H$ is an isometry we have $\mathfrak{s}_1(\mathbf{V}) = \mathfrak{s}_1(\nu)$ and $\mathfrak{s}_{Nm}(\mathbf{V}) = \mathfrak{s}_N(\nu)$. Rewrite the distribution as

$$\begin{aligned}\mathcal{D}^{\tau}_{R^m, \sqrt{\mathbf{\Sigma}}, \mathbf{c}} := \mathcal{D}_{\tau(R^m), \sqrt{\mathbf{\Sigma}}, \tau(\mathbf{c})} &= \mathbf{V}^{-1} \cdot \mathcal{D}_{\mathbf{V} \cdot \tau(R^m), \sqrt{\mathbf{V\Sigma V}^T}, \mathbf{V}\tau(\mathbf{c})} \\ &= \mathbf{V}^{-1} \cdot \mathcal{D}_{\theta(R^m), \sqrt{\mathbf{V\Sigma V}^T}, \theta(\mathbf{c})} =: \mathbf{V}^{-1} \cdot \mathcal{D}^{\theta}_{R^m, \sqrt{\mathbf{\Sigma}}, \mathbf{c}},\end{aligned}$$

which holds by Lemma 14 since $\mathbf{V}$ is injective. Besides, for any $\mathbf{x} \in \theta(R^m)$

$$\begin{aligned}(\mathbf{V}^{-1} \cdot \mathbf{x}) \bmod \tau(\Lambda_q^{\perp}(\mathbf{A})) &= \mathbf{V}^{-1} \cdot \mathbf{x} \bmod \mathbf{V}^{-1} \cdot \theta(\Lambda_q^{\perp}(\mathbf{A})) \\ &= \mathbf{V}^{-1} \cdot \mathbf{x} + \mathbf{V}^{-1} \cdot \theta(\Lambda_q^{\perp}(\mathbf{A})) \\ &= \mathbf{V}^{-1} \cdot (\mathbf{x} \bmod \theta(\Lambda_q^{\perp}(\mathbf{A}))).\end{aligned}$$

One of the conditions in Theorem 6 states $\Pr\big[\eta_{\varepsilon}(\theta(\Lambda_q^{\perp}(\mathbf{A}))) < \beta | \mathbf{A} \leftarrow \mathcal{P}_A\big] \geq 1 - \delta_1$. Say for a fixed matrix $\mathbf{A} \in R_q^{n \times m}$ this condition is true. Then since we require $s_{Nm}(\mathbf{V})^2 \cdot \mathbf{\Sigma} \geq \beta^2$, we have $\mathbf{V\Sigma V}^T \geq \eta_{\varepsilon}(\theta(\Lambda_q^{\perp}(\mathbf{A})))^2$. For such matrix we apply Corollary 1 as follows

$$\begin{aligned}\mathcal{D}^{\tau}_{R^m, \sqrt{\mathbf{\Sigma}}, \mathbf{c}} \bmod \tau(\Lambda_q^{\perp}(\mathbf{A})) &= \mathbf{V}^{-1} \cdot \Big(\mathcal{D}^{\theta}_{R^m, \sqrt{\mathbf{V\Sigma V}^T}, \mathbf{c}} \bmod \theta(\Lambda_q^{\perp}(\mathbf{A}))\Big) \\ &\approx_{\varepsilon/(1-\varepsilon)} \mathbf{V}^{-1} \cdot \mathcal{U}(\theta(R^m)/\theta(\Lambda_q^{\perp}(\mathbf{A}))) \\ &= \mathcal{U}(\tau(R^m)/\tau(\Lambda_q^{\perp}(\mathbf{A}))).\end{aligned}$$

Thus $\mathrm{SD}(\mathcal{D}^{\theta}_{R^m,\sqrt{\boldsymbol{\Sigma}},\mathbf{c}} \bmod \tau(\Lambda_q^{\perp}(\mathbf{A})), \mathcal{U}(\tau(R^m)/\tau(\Lambda_q^{\perp}(\mathbf{A})))) \leq \varepsilon/(1-\varepsilon)$ holds for any $\boldsymbol{\Sigma} \geq \beta^2/\mathfrak{s}_N(\nu)^2$. Taking into account the remaining conditions of Theorem 6 independent of the embedding for a random $\mathbf{A} \leftarrow \mathcal{P}_A$ we obtain

$$\mathrm{SD}((\mathbf{Ax} \bmod q, \mathbf{A}), (\mathbf{u}, \mathbf{A})) \leq \varepsilon/(1-\varepsilon) + \delta_1 + \delta_2.$$

The statements for the Rényi divergence and the implication from coefficient embedding to canonical embedding are obtained in the same way. □

4.2 Smoothing Analysis

The quality of Theorem 6 heavily depends on the shape of the distribution $\mathcal{P}_A$, as well as the bound β on the smoothing parameter, determining the probabilities δ_1 and δ_2. The following results give bounds on β and δ_1 for $\mathcal{P}_A$ equal to the uniform distribution of $n \times m$ matrices over R_q. Note that for this setting, Lemma 21 provides bounds on the probability δ_2, i.e. that the image of $\mathbf{A}$ has full rank.

Overall, we distinguish two different strategies to analyse the smoothing parameter of $\Lambda_q^{\perp}(\mathbf{A})$. The first strategy was given in [LPR13], it directly computes the expected Gaussian weight of the dual lattice. The second strategy was detailed in [SS11] and its follow-up work [RSW18]. It focuses on lower bounding the length of the shortest vector in the dual lattice.

We first present statements that appeared in prior work where the distribution of $\mathbf{A}$ is not truly uniform. Then we give generalisations of both strategies to matrices uniform in $R_q^{n\times m}$, we showcase the proofs in both canonical and coefficient embedding even though both can be applied to an arbitrary embedding.

The first classic result comes from [SS11] and considers a uniform matrix with invertible entries.

Lemma 22 ([SS11, Theorem 2]). *Let K be a cyclotomic number field of degree N with R its ring of integers, where $N \geq 8$ is a power of 2. Let $q \geq 5$ be a prime, $m \geq 2$ be an integer, $\varepsilon \in (0, 1/2)$ and $\varepsilon' > 0$.*

Let $r = \sqrt{\ln(2Nm(1+1/\varepsilon))/\pi} \cdot \min(\sqrt{N} \cdot q^{\frac{1}{m}+\varepsilon'}, q^{\frac{1}{m}+f\cdot\varepsilon'})$. Then

$$\Pr[\eta_{\varepsilon}(\tau(\Lambda_q^{\perp}(\mathbf{a}))) > r \mid \mathbf{a} \leftarrow \mathcal{U}((R_q^{\times})^m)] \leq 2^{3N(m+1)} \cdot q^{-\varepsilon' m N}.$$

Its canonical embedding version is described in [RSW18].

Lemma 23 ([RSW18, Lemma 5.2]). *Let K be a cyclotomic number field of degree N with R its ring of integers. Let $q \geq 2$ be a prime not dividing Δ_K. Let $m \geq 2$, $\varepsilon \in (0, 1/2)$ and $\varepsilon' > 0$.*

Then for $r = \sqrt{\ln(2Nm(1+1/\varepsilon))/\pi} \cdot \Delta_K^{1/N} \cdot q^{\frac{1}{m}+\varepsilon'}$ it holds

$$\Pr[\eta_{\varepsilon}(\theta(\Lambda_q^{\perp}(\mathbf{a}))) > r \mid \mathbf{a} \leftarrow \mathcal{U}((R_q^{\times})^m)] \leq 2^{3N(m+1)} \cdot q^{-\varepsilon' m N}.$$

We also state the original [LPR13] statement that covers uniform matrices with prepended identity.

Lemma 24 ([LPR13, Theorem 7.4]). *Let K be a cyclotomic number field of degree N with R its ring of integers. Let $q > 2$, $0 < n \leq m$ be integers, $\varepsilon = q^2 \cdot m \cdot 2^{1-2N}$ and $s > 0$ s.t. $2^{2s} > 2 \cdot (2^{-N(m-2)} + q^2 \cdot m)$. Let $\mathcal{P}_A = \left\{\mathbf{A} = [\mathbf{I}_n \mid \bar{\mathbf{A}}] \mid \bar{\mathbf{A}} \leftarrow \mathcal{U}\left(R_q^{n\times(m-n)}\right)\right\}$. Then for every $r > 2N$*

$$\mathbb{E}_{\mathbf{A}\leftarrow\mathcal{P}_A}\left[\rho_{1/r}(\theta(\Lambda_q^{\perp}(\mathbf{A}))^*)\right] \leq 1 + 2(r/N)^{-Nm} \cdot q^{Nn+2} + \varepsilon.$$

If $r > 2N \cdot q^{n/m+2/Nm}$ then

$$\Pr\left[\eta_{2^{-N+s}}(\theta(\Lambda_q^{\perp}(\mathbf{A}))) > r \mid \mathbf{A} \leftarrow \mathcal{P}_A\right] \leq 2^{-N+s}.$$

The following lemma can be seen as a generalisation of their results to completely uniform matrices. Unfortunately, considering this less structured distribution makes the parameters depend on the factorisation of the integer modulus q. Contrary to the original work we restrict q to be a prime, we delegate lifting this restriction to future work.

Lemma 25 ([LPR13] for Uniform Matrices). *Let K be a cyclotomic field of degree N with R its ring of integers. Further, let $q > 2$ be an unramified prime modulus such that $R_q = \prod_{j=1}^{f} \mathbb{Z}_q[X]/\langle \Phi_j(X) \rangle$ for distinct, irreducible $\Phi_j(X)$ of degree δ. Let $0 < n \leq m, 2 \leq m$ be integers and $s > 0$, $\varepsilon = m \cdot 2^{1-2N} + 2f \cdot q^{-\delta(m-n)} \cdot (1 + m \cdot 2^{1-2N})$ such that $\delta(m-n)\log q > 2N$, $2^{2s} > 2 \cdot (2^{-N(m-2)} + m + f \cdot (1 + m \cdot 2^{1-2N}))$, $2f \cdot q^{-\delta(m-n)} \leq 1$. Then for every $r > 2N$*

$$\mathbb{E}_{\mathbf{A} \leftarrow \mathcal{U}(R_q^{n\times m})} \left[\rho_{1/r}(\theta(\Lambda_q^{\perp}(\mathbf{A}))^*)\right] \leq 2\exp(1) \cdot (r/N)^{-Nm} \cdot q^{Nn} + 1 + \varepsilon.$$

If $r > 2\exp(1/Nm) \cdot N \cdot q^{n/m}$ then

$$\Pr\left[\eta_{2^{-N+s}}(\theta(\Lambda_q^{\perp}(\mathbf{A}))) > r \mid \mathbf{A} \leftarrow \mathcal{U}(R_q^{n\times m})\right] \leq 2^{-N+s}.$$

Proof. First, we identify the exact shape of the dual lattice of $\Lambda_q^{\perp}(\mathbf{A})$ from Lemma 20 for the canonical embedding

$$\theta(\Lambda_q^{\perp}(\mathbf{A}))^* = \theta\left((R^*)^m + \left\{\frac{1}{q}\mathbf{A}^T \cdot \mathbf{s} \mid \mathbf{s} \in (R^*)^n\right\}\right).$$

For cleaner notations we drop mentioning the embedding $\theta(\cdot)$ in the rest of the proof and denote the expected value we analyse as $E \coloneqq \mathbb{E}_{\mathbf{A} \leftarrow \mathcal{U}(R_q^{n\times m})}\left[\rho_{1/r}(\Lambda_q^{\perp}(\mathbf{A})^*)\right]$. For arbitrary rings $\{R_i\}_{i=0,1,2}$ where $R_2 = R_0/R_1$ is a quotient of R_0 we define a function rep : $R_2 \to R_0$ that maps an element of R_2 to its arbitrary representative in R_0. For the expectation we have

$$\begin{aligned} E &\leq \mathbb{E}_{\mathbf{A} \leftarrow \mathcal{U}(R_q^{n\times m})}\left[\sum_{\mathbf{s}\in(R_q^*)^n} \rho_{1/r}\left(1/q \cdot ((qR^*)^m + \mathrm{rep}(\mathbf{A}^T \cdot \mathbf{s}))\right)\right] \\ &= \sum_{\mathbf{s}\in(R_q^*)^n} \mathbb{E}_{\mathbf{A} \leftarrow \mathcal{U}(R_q^{n\times m})}\left[\rho_{1/r}\left(1/q \cdot ((qR^*)^m + \mathrm{rep}(\mathbf{A}^T \cdot \mathbf{s}))\right)\right] \\ &= \sum_{\mathbf{s}\in(R_q^*)^n} \left(\mathbb{E}_{\mathbf{a} \leftarrow \mathcal{U}(R_q^{n})}\left[\rho_{1/r}\left(1/q \cdot (qR^* + \mathrm{rep}(\mathbf{a}^T \cdot \mathbf{s}))\right)\right]\right)^m, \end{aligned}$$

where the first equality holds by linearity of the expectation and the second holds since the rows of $\mathbf{A}$ are iid. Note that since we add qR^* to a representative of $\mathbf{a}^T \cdot \mathbf{s} \in R^*/qR^*$ the value above is independent of the specific choice of the representative. By Lemma 10, the value $\mathbf{a}^T \cdot \mathbf{s} \bmod q$ is distributed uniformly in the ideal $\mathcal{I}_{\mathbf{s}} \subset R_q^*$. Then

$$\begin{aligned} \mathbb{E}\left[\rho_{1/r}\left(1/q \cdot (qR^* + \mathrm{rep}(\mathbf{a}^T \cdot \mathbf{s}))\right)\right] &= \sum_{x\in\mathcal{I}_{\mathbf{s}}} \frac{1}{|\mathcal{I}_{\mathbf{s}}|} \cdot \rho_{1/r}\left(1/q \cdot (qR^* + \mathrm{rep}(x))\right) \\ &= \frac{1}{|\mathcal{I}_{\mathbf{s}}|} \cdot \rho_{1/r}\left(1/q \cdot (qR^* + \mathrm{rep}(\mathcal{I}_{\mathbf{s}}))\right). \end{aligned}$$

Now we group the summands in $\mathbb{E}_{\mathbf{A} \leftarrow \mathcal{U}(R_q^{n\times m})}\left[\rho_{1/r}(\Lambda_q^{\perp}(\mathbf{A})^*)\right]$ based on the ideal $\mathbf{s}$ generates. Denote $\mathcal{I}' = qR^* + \mathrm{rep}(\mathcal{I}) \subset R^*$ an ideal of R^* containing qR^* that corresponds to an ideal $\mathcal{I} \subset R_q^*$. Following Lemma 5 we have another bijection $R = t \cdot R^*$ for a ring element t. It extends to the ideals of R^* denoting $\mathcal{J}' := t \cdot \mathcal{I}' = t \cdot \mathrm{rep}(\mathcal{I}) + qR$ and to the ideals of R_q^* denoting $\mathcal{J} := t \cdot \mathcal{I} \subset R_q$. Then $\mathcal{I} = \mathcal{I}' \bmod qR$, $\mathcal{J} = \mathcal{J}' \bmod qR$ and $|\mathcal{J}| = |\mathcal{I}|$. Then we have

$$\begin{aligned} N(1/q \cdot \mathcal{I}') = \frac{N(\mathcal{I}')}{|N(q)|} &= \frac{N(\mathcal{J}')}{|N(t)| \cdot |N(q)|} = \frac{N(\mathcal{J})}{|N(t)| \cdot |N(q)|} \\ &= \frac{|R_q|}{|\mathcal{J}| \cdot \Delta_K \cdot |R_q|} = \frac{1}{|\mathcal{J}| \cdot \Delta_K}. \end{aligned} \tag{4}$$

In the transitions above we use multiplicativity of norms, the fact that for $\mathcal{J} = \mathcal{J}' \bmod qR$ we have $N(\mathcal{J}) = N(\mathcal{J}')$ and the definition of the norm in finite rings. Then if T is the set of all ideals $\mathcal{I}' \subset R^*$ that contain qR^* using the notation above we have

$$\begin{aligned} E &\leq \sum_{\mathcal{I}' \in T} \left(\frac{1}{|\mathcal{I}' \bmod qR^*|} \cdot \rho_{1/r}(1/q \cdot \mathcal{I}') \right)^m \sum_{\mathbf{s}: \mathcal{I}_{\mathbf{s}} = \mathcal{I}' \bmod qR^*} 1 \\ &\leq \sum_{\mathcal{I}' \in T} \left(\frac{1}{|\mathcal{I}' \bmod qR^*|} \cdot \rho_{1/r}(1/q \cdot \mathcal{I}') \right)^m \cdot |\mathcal{I}' \bmod qR^*|^n \\ &= \sum_{\mathcal{I}' \in T} \frac{1}{|\mathcal{I}|^{m-n}} \cdot \rho_{1/r}(1/q \cdot \mathcal{I}')^m = \sum_{\mathcal{I}' \in T} \frac{1}{|\mathcal{J}|^{m-n}} \cdot \rho_{1/r}(1/q \cdot \mathcal{I}')^m. \end{aligned}$$

Now for each ideal $\mathcal{I}' \in T$ by Lemma 19 and Equation (4)

$$\begin{aligned} \rho_{1/r}(1/q \cdot \mathcal{I}')^m &\leq \max(1, r^{-Nm} \cdot N(1/q \cdot \mathcal{I}')^{-m})(1 + 2^{-2N})^m \\ &\leq \max(1, r^{-Nm} \cdot (|\mathcal{J}| \cdot \Delta_K)^m)(1 + 2^{-2N})^m \\ &\leq \underbrace{1 + m2^{1-2N}}_{\alpha} + \underbrace{2r^{-Nm}\Delta_K^m}_{\beta} \cdot |\mathcal{J}|^m. \end{aligned}$$

where $\mathcal{J} \subset R_q$ is defined as above. In the last inequality $(1+2^{-2N})^m \leq \frac{1}{1-m \cdot 2^{-2N}} \leq 1+2 \cdot m \cdot 2^{-2N}$ since $m \cdot 2^{-2N} \leq 1/2$. For the case when $r^{-Nm} \cdot (|t \cdot \mathcal{J} \bmod qR| \cdot \Delta_K)^m > 1$ we use $\frac{1}{1-m \cdot 2^{-2N}} \leq 2$. To bound the larger of two values we take the sum of the bounds. Denote $\mathbb{I}$ the set of ideals in R_q. Putting the arguments together we get

$$\begin{aligned} E &\leq \alpha \cdot \sum_{\mathcal{J} \in \mathbb{I}} |\mathcal{J}|^{n-m} + \beta \cdot \sum_{\mathcal{J} \in \mathbb{I}} |\mathcal{J}|^n \\ &\leq \alpha \cdot (1 + 2fq^{-\delta(m-n)}) + 2\exp(1) \cdot r^{-Nm} N^{Nm} q^{Nn} \end{aligned}$$

where in the last inequality we use Lemma 6. Lastly we apply the Markov inequality. Define arbitrary positive values $\varepsilon', \varepsilon''$ such that $\varepsilon' \cdot \varepsilon'' > \varepsilon$. Then for any $r > \exp(1/Nm) \cdot N \cdot q^{n/m} \cdot (2/(\varepsilon'' \cdot \varepsilon' - \varepsilon))^{1/Nm}$ we have

$$\begin{aligned} \Pr(\rho_{1/r}(\Lambda_q^{\perp}(\mathbf{A})^*) - 1 > \varepsilon') &\leq \frac{\mathbb{E}_{\mathbf{A}}\left[\rho_{1/r}(\Lambda_q^{\perp}(\mathbf{A})^*)\right] - 1}{\varepsilon'} \\ &\leq \frac{m2^{1-2N} + 2fq^{-\delta(m-n)} \cdot (1 + m2^{1-2N}) + \varepsilon''\varepsilon' - \varepsilon}{\varepsilon'} \\ &\leq \varepsilon''. \end{aligned}$$

Here in the last inequality we used the definition of ε from the lemma statement. Set $\varepsilon' = \varepsilon'' = 2^{-N+s}$. Then for $\delta(m-n)\log q > 2N$ and $2^{2s} > 2 \cdot (2^{-N(m-2)} + m + f \cdot (1 + m \cdot 2^{1-2N}))$ we have $\varepsilon' \cdot \varepsilon'' > \varepsilon$ and $\left(\frac{2}{\varepsilon'' \cdot \varepsilon' - \varepsilon}\right)^{1/Nm} \leq 2$, and the statement follows. □

In the original result of [SS11,RSW18] , the distribution $\mathcal{P}_A$ enforced a full image rank by letting every entry of $\mathbf{A}$ be invertible in R_q. The following result can be seen as a generalisation to the setting where $\mathcal{P}_A$ is the uniform distribution over $R_q^{n\times m}$, proven in the coefficient embedding.

Lemma 26 ([SS11,RSW18] for Uniform Matrices). *Let K be a cyclotomic field of degree N with R its ring of integers. Further, let $q \geq 5$ be an unramified prime modulus such that $R_q = \prod_{j=1}^{f} \mathbb{Z}_q[X]/\langle \Phi_j(X)\rangle$ for distinct, irreducible $\Phi_j(X)$ of degree δ. Let $N \geq 8$ be a power of 2. Let $\max(n,2) \leq m$ be integers, $\varepsilon \in (0,1/2)$ and $\varepsilon' > 0$.*

Let $r = \sqrt{\ln(2Nm(1+1/\varepsilon))/\pi} \cdot \sqrt{N} \cdot q^{n/m} \cdot e^{\varepsilon'}$. Then

$$\Pr[\eta_\varepsilon(\tau(\Lambda_q^\perp(\mathbf{A}))) > r \mid \mathbf{A} \leftarrow \mathcal{U}(R_q^{n\times m})] \leq 2^{-Nm(\varepsilon'-3)}$$

Proof. Following Lemma 20 the dual lattice to $\Lambda_q^\perp(\mathbf{A})$ with respect to the coefficient embedding is of the form

$$\begin{aligned}\tau(\Lambda_q^\perp(\mathbf{A}))^* &= \frac{1}{q}\Lambda_q(\mathrm{Rot}(\mathbf{A})^T)\\ &= \frac{1}{q}\big\{\tau(\mathbf{t}) \mid \mathbf{t} \in R^m \text{ s.t. } \exists \mathbf{s} \in R_q^n : \tau(\mathbf{t}) = \mathrm{Rot}(\mathbf{A})^T \cdot \tau(\mathbf{s}) \bmod q\big\}\\ &= \frac{1}{q}\tau(\Lambda_q(\bar{\mathbf{A}}^T)).\end{aligned}$$

By properties of power-of-2 cyclotomics, the matrix $\mathrm{Rot}(\mathbf{A})^T$ is in bijective correspondence with $\mathrm{Rot}(\bar{\mathbf{A}}^T)$ for some $\bar{\mathbf{A}} \in R_q^{n\times m}$. The bijection maps every coordinate $a_{ij} = \sum_{k=0}^{N-1} \alpha_k X^k \in R_q$ of $\mathbf{A}$ to a corresponding coordinate of $\bar{\mathbf{A}}$ equal $\bar{a}_{ij} = \alpha_0 - \sum_{k=1}^{N-1} \alpha_k X^{N-k}$. Therefore, $\bar{\mathbf{A}}$ follows the uniform distribution and it suffices for us to lower bound the first minimum of the lattice $\lambda_1^\infty(\frac{1}{q}\tau(\Lambda_q(\mathbf{A}^T)))$ for $\mathbf{A} \leftarrow \mathcal{U}(R_q^{n\times m})$.

We claim that $\lambda_1^\infty(\frac{1}{q}\tau(\Lambda_q(\mathbf{A}^T))) \geq 1/(\sqrt{N}\cdot q^{\varepsilon'/\log q + n/m})$ except with probability $\leq 2^{-Nm(\varepsilon'-3)}$. Let $B > 0$ and then apply Lemma 11. For every $\mathbf{t} \in R^m$ of norm $0 < \|\mathbf{t}\|_\infty < B$ and $\mathbf{s} \in R_q^n$ define

$$p(\mathbf{t},\mathbf{s}) = \Pr_{\{\mathbf{a}_i\}_i \leftarrow (R_q^n)^m}(\forall i : t_i = \mathbf{a}_i^T \cdot \mathbf{s} \bmod q) = \prod_{i=1}^{m} \Pr_{\mathbf{a}_i \leftarrow R_q^n}(t_i = \mathbf{a}_i^T \cdot \mathbf{s} \bmod q).$$

Denote $p_i(t_i,\mathbf{s}) := \Pr_{\mathbf{a}_i \leftarrow R_q^n}(t_i = \mathbf{a}_i^T \cdot \mathbf{s} \bmod q)$. By Lemma 10 the value $\mathbf{a}_i^T \cdot \mathbf{s} \bmod q$ is distributed uniformly in $\mathcal{I}_\mathbf{s}$. Then for any fixed $t_i \in \mathcal{I}_\mathbf{s}$

$$p_i(t_i,\mathbf{s}) = \frac{1}{|\mathcal{I}_\mathbf{s}|}.$$

And the probability p that the lattice contains a short infinity norm vector is bounded as

$$\begin{aligned}p &= \Pr_{\mathbf{A}}(\exists \mathbf{t} : \mathbf{t} \in \tau(\Lambda_q(\mathbf{A}^T)), 0 < \|\mathbf{t}\|_\infty < B)\\ &\leq \sum_{0<\|\mathbf{t}\|_\infty<B, \mathbf{s}\in R_q^n} p(\mathbf{t},\mathbf{s})\\ &= \sum_{0<\|\mathbf{t}\|_\infty<B, \mathbf{s}\in R_q^n} \prod_i p_i(\mathbf{t}_i,\mathbf{s})\\ &\leq \sum_{k=0}^{f} \sum_{\substack{\mathcal{I}\in\mathbb{I}\\ |\mathcal{I}|=q^{\delta k}}} \sum_{\mathbf{s}\in\mathcal{I}^n} \sum_{\substack{0<\|\mathbf{t}\|_\infty<B\\ \mathbf{t}\in\mathcal{I}^m}} \frac{1}{q^{m\cdot\delta k}}\\ &\leq |\{\mathcal{I}\in\mathbb{I}\}| \cdot \max_{k\in[f]} q^{n\cdot\delta k} \cdot \frac{N(B,k)^m}{q^{m\cdot\delta k}} = 2^f \cdot \max_{k\in[f]} \frac{N(B,k)^m}{q^{(m-n)\cdot\delta k}}\end{aligned}$$

The first inequality comes from the union bound. The second inequality conditions on the ideal generated by coefficients of $\mathbf{s}$. We enumerate over all ideals of R_q according to their norm. We also substitute the generating condition by $\mathbf{s} \in \mathcal{I}^n$. In the third inequality $N(B,k)$ denotes the number of elements $t \in R_q$ such that $0 < \|t\|_\infty < B, t \in \mathcal{I}$ for $\mathcal{I}$ of cardinality $q^{\delta k}$. Then $|\mathbf{s} \in \mathcal{I}^n| = q^{n \cdot \delta k}$. We upper bound this probability over all ideals of R_q.

We now analyse $N(B,k)$. For convenience, let $S = \{j_1, \ldots, j_{f-k}\}$ be an arbitrary subset of $\{1, \ldots, f\}$ of cardinality $f-k$. The bound we obtain in the end only depends on the cardinality of S. We proceed by a volume argument.

Denote $\lambda_\infty = \lambda_1^\infty(\tau(\mathcal{I}))$ for $\mathcal{I} = \langle q, \prod_{j \in S} \Phi_j(X) \rangle$. Then $N(B,k)$ is upper bounded by the ratio between volumes of hypercubes $\left\{\mathbf{v} \in \mathbb{R}^N \mid \|\mathbf{v}\|_\infty \leq B + \lambda_\infty/2\right\}$ and $\left\{\mathbf{v} \in \mathbb{R}^N \mid \|\mathbf{v}\|_\infty \leq \lambda_\infty\right\}$

$$N(B,k) \leq \left(\frac{2B}{\lambda_\infty} + 1\right)^N$$

By Lemma 17 and inequalitites between norms $\frac{1}{\mathfrak{s}_1(\nu)} N(\mathcal{I})^{1/N} \leq \frac{1}{\sqrt{N}} \cdot \lambda_1(\tau(\mathcal{I})) \leq \lambda_\infty$. For the power-of-2 cyclotomics $\mathfrak{s}_1(\nu) = \sqrt{N}$. Computing the norm we get $1/\sqrt{N} \cdot q^{\delta(f-k)/N} = 1/\sqrt{N} \cdot q^{1-k/f} \leq \lambda_\infty$. Take $B = \frac{1}{\sqrt{N}} q^\beta$ then for $\beta < 1 - k/f$ we have $\lambda_\infty > B$ and $N(B,k) = 0$. Otherwise, for large enough q we have $2q^{\beta - (1-k/f)} > 1$ and

$$N(B,k) \leq (2q^{\beta - 1 + k/f} + 1)^N \leq 2^{2N} \cdot q^{N\beta - N + \delta k}.$$

Here in the first expression we insert the computed bounds and in the second inequality we use $(x+1)^N \leq (2x)^N$ for $x > 1$. Hence,

$$\begin{aligned} p &\leq 2^f \cdot \max_{k \in [f]} \frac{2^{2Nm} q^{Nm\beta - Nm + \delta km}}{q^{(m-n) \cdot \delta k}} \\ &\leq 2^{N(2m+1)} q^{Nm\beta - Nm + Nn} \leq 2^{-Nm(\varepsilon' - 3)}. \end{aligned}$$

In the first inequality use that the maximum is attained for $k = f$ and the inequality $f < N$ for the other term. In the last equality we set $\beta = 1 - n/m - \varepsilon'/\log q$ and bound $2^{N(2m+1)} \leq 2^{3Nm}$. It follows that the shortest vector of $1/q \cdot \tau(\Lambda_q(\mathbf{A}^T))$ is at least $\frac{1}{\sqrt{N}} q^{\beta - 1} = \frac{1}{\sqrt{N}} \cdot q^{-n/m - \varepsilon'/\log q}$ with probability in $1 - 2^{-Nm(\varepsilon' - 3)}$. Using the conventional smoothing bound from Lemma 11 and $q^{1/\log q} = e$ we obtain the result. □

4.3 Smoothing in the Presence of Exact Linear Leakage

In the following, we study the smoothing-based approach under linear leakage on the hash input vector $\mathbf{x}$. We distinguish between *exact* and *noisy* leakage. The first type of leakage is addressed in this section and the proof technique is inspired by a recent LHL variant [SSE+24], lifting it to a more general leakage setting. The second class of leakage is treated in Section 4.4, using recent results [ENP24] on the distribution of discrete Gaussians under noisy linear leakage. It can be seen as a generalization of the first type of leakage studied in [DGKS21].

In the exact case we consider leakage of the form $(\mathbf{E}, \mathbf{E} \cdot \mathbf{x})$ for a given wide matrix $\mathbf{E} = (\mathbf{I}_l \mid \bar{\mathbf{E}}) \in R^{l \times m}$ with $l < m$. The matrix $\bar{\mathbf{E}}$ can be sampled from a short distribution or be fixed as long as it satisfies the required conditions.

We rely heavily on the approach of [SSE+24]. The original work defined matrices $\mathbf{A}$ and $\mathbf{E}$ so that the leakage is not exact and corresponds to our Section 4.4. However, their techniques are powerful enough to cover the case of exact leakage that we present here. The original work also covers composite modulus q which we omit for simplicity.

Theorem 7. *Let K be a cyclotomic number field with N a power of 2. Let q be a prime, $0 < n + l \leq m$ be integers. Let $\varepsilon \in (0, 1/2)$, $\delta_E, \delta_2 \in (0, 1)$, $\varepsilon' > 3$, $B_{\max} > 0$, and let $\mathbf{\Sigma} \in \mathbb{R}^{Nm \times Nm}$ be positive-definite. Assume*

$$\mathbf{\Sigma} \geq Nm \cdot \max\left(2B_{\max}, \sqrt{N}e^{\varepsilon'} \cdot q^{(n+l)/m}\right) \cdot \sqrt{\ln(2N(m-l)(1+1/\varepsilon))/\pi},$$
$$q \geq 2 \cdot B_{\max} \cdot Nm \cdot \max\left(2B_{\max}, \sqrt{N}e^{\varepsilon'} \cdot q^{(n+l)/m}\right).$$

Let $\mathcal{P}_E$ be a distribution on the set of matrices $R^{l \times m}$ s.t. with probability at least $1 - \delta_E$ for $\mathbf{E} = [\mathbf{I}_l \mid \bar{\mathbf{E}}] \leftarrow \mathcal{P}_E$ it holds

1. $\mathfrak{s}_1(\mathrm{Rot}(\mathbf{E})) < B_{\max}$,
2. $\Pr\left[\mathbf{A} \cdot \Lambda^\perp(\mathbf{E}) \bmod q = R_q^n \mid \mathbf{A} \leftarrow \mathcal{U}(R_q^{n \times m})\right] \geq 1 - \delta_2$.

Then for distributions defined as

$$\mathcal{P}' = \left\{\mathbf{Ax} \bmod q, \mathbf{Ex}, \mathbf{A}, \mathbf{E} \mid \mathbf{x} \leftarrow \mathcal{D}^\tau_{R^m, \mathbf{\Sigma}}\right\},$$
$$\mathcal{U}' = \left\{\mathbf{u}, \mathbf{Ex}, \mathbf{A}, \mathbf{E} \mid \mathbf{x} \leftarrow \mathcal{D}^\tau_{R^m, \mathbf{\Sigma}}, \mathbf{u} \leftarrow \mathcal{U}(R_q^n)\right\},$$

we have

$$\mathrm{SD}(\mathcal{P}', \mathcal{U}') \leq \varepsilon/(1-\varepsilon) + 2^{-Nm(\varepsilon'-3)} + \delta_E + \delta_2,$$
$$\mathrm{RD}(\mathcal{P}'; \mathcal{U}') < \frac{1+\varepsilon}{1-\varepsilon} + (2^{-Nm(\varepsilon'-3)} + \delta_E + \delta_2) \cdot q^{Nn}.$$

Proof. Similarly to Theorem 6 let $\mathcal{S}$ represent the set of all matrices $(\mathbf{A}, \mathbf{E})$ for which conditions (1, 2) hold. Consider an arbitrary pair $(\mathbf{A}, \mathbf{E}) \in \mathcal{S}$.

For any leakage $\mathbf{w} \in \mathrm{im}(\mathbf{E})$ ($= \mathbf{E} \cdot \mathbf{t}'$ for some $\mathbf{t}'$) there exists a vector $\mathbf{t} \in R^m$ s.t. $\mathbf{A} \cdot \mathbf{t} = \mathbf{0} \bmod q$ and $\mathbf{E} \cdot \mathbf{t} = \mathbf{w}$. We construct it as $\mathbf{t} = \mathbf{t}' + \mathbf{t}''$ where we denote $\mathbf{r} := \mathbf{At}' \bmod q$ and find $\mathbf{t}'' \in \Lambda^\perp(\mathbf{E})$ s.t. $\mathbf{A} \cdot \mathbf{t}'' = -\mathbf{r} \bmod q$. Such $\mathbf{t}''$ exists by Condition 2 of the theorem. Then given the leakage $\mathbf{w}$ the distribution of $\mathbf{x}$ is

$$\left\{\mathbf{x} \leftarrow \mathcal{D}_{R^m, \mathbf{\Sigma}} \mid \mathbf{x} \in (\Lambda^\perp(\mathbf{E}) + \mathbf{t})\right\} = \mathcal{D}_{\Lambda^\perp(\mathbf{E}) + \mathbf{t}, \mathbf{\Sigma}} = \mathbf{t} + \mathcal{D}_{\Lambda^\perp(\mathbf{E}), \mathbf{\Sigma}, -\mathbf{t}},$$

and it is enough to analyse $\mathbf{A} \cdot (\mathbf{t} + \mathcal{D}_{\Lambda^\perp(\mathbf{E}), \mathbf{\Sigma}, -\mathbf{t}}) = \mathbf{A} \cdot \mathcal{D}_{\Lambda^\perp(\mathbf{E}), \mathbf{\Sigma}, -\mathbf{t}} \bmod q$. The matrix $\mathbf{A} : \Lambda^\perp(\mathbf{E}) \to R_q^n$ represents a linear map. It generates an isomorphism $\Lambda^\perp(\mathbf{E}) \bmod (\Lambda_q^\perp(\mathbf{A}) \cap \Lambda^\perp(\mathbf{E})) \sim \mathbf{A} \cdot \Lambda^\perp(\mathbf{E}) \bmod q = R_q^n$. Here the last equality follows again from Condition 2. Hence, it remains to prove that $\mathcal{D}_{\Lambda^\perp(\mathbf{E}), \mathbf{\Sigma}, -\mathbf{t}} \bmod (\Lambda_q^\perp(\mathbf{A}) \cap \Lambda^\perp(\mathbf{E}))$ is uniformly distributed. By Corollary 1 and Corollary 2 the distribution is close to uniform if $\mathbf{\Sigma} \geq \eta_\varepsilon(\tau(\Lambda_q^\perp(\mathbf{A}) \cap \Lambda^\perp(\mathbf{E})))$. In the remainder of the proof we analyse the smoothing parameter of this lattice.

Since $\mathbf{E} = [\mathbf{I} \mid \bar{\mathbf{E}}]$ contains the identity, the matrix $\mathrm{Rot}(\mathbf{E})$ is full rank hence the lattice $\tau(\Lambda^\perp(\mathbf{E}))$ is of rank $N(m-l)$ in $\mathbb{R}^{Nm}$. We know that $\tau(\Lambda_q^\perp(\mathbf{A}))$ is full rank in the same dimension. Hence, their intersection is a lattice with the same rank as $\tau(\Lambda^\perp(\mathbf{E}))$. It remains to upper bound $\lambda_{N(m-l)}(\tau(\Lambda_q^\perp(\mathbf{A}) \cap \Lambda^\perp(\mathbf{E})))$ and apply the classic smoothing parameter bound in Lemma 12.

Consider a simpler lattice $\Lambda = \Lambda_q^\perp(\mathbf{A}) \cap \Lambda_q^\perp(\mathbf{E})$ of rank Nm where the kernel of $\mathbf{E}$ is taken modulo q. By the transference theorems

$$\lambda_{N(m-l)}(\tau(\Lambda)) \leq \frac{Nm}{\lambda_{Nl+1}(\tau(\Lambda)^*)} = \frac{Nm}{1/q \cdot \lambda_{Nl+1}(\Lambda_q([\mathrm{Rot}(\mathbf{A})^T | \, \mathrm{Rot}(\mathbf{E})^T]))},$$

where we know the shape of the dual lattice by Lemma 20. Denote $\mathbf{M} = [\mathrm{Rot}(\mathbf{A}) || \, \mathrm{Rot}(\mathbf{E})]$. Since $\mathrm{Rot}(\mathbf{E})^T$ is only of rank Nl, the $(Nl+1)$st minimum of $\Lambda_q(\mathbf{M}^T)$ has to be larger than some vector outside the span of $\mathrm{Rot}(\mathbf{E})^T$. Formally, $\lambda_{Nl+1}(\Lambda_q(\mathbf{M}^T)) \geq \lambda_1(\Lambda_q(\mathbf{M}^T) \setminus \mathrm{Rot}(\mathbf{E})^T \cdot \mathbb{Z}^{Nl})$.

We represent $\tau(\mathbf{u}) \in \Lambda_q(\mathbf{M}^T) \setminus \mathrm{Rot}(\mathbf{E})^T \cdot \mathbb{Z}^{Nl}$ as $\mathbf{u} = \mathrm{Rot}(\mathbf{A})^T \cdot \tau(\mathbf{v}_A) + \mathrm{Rot}(\mathbf{E})^T \cdot \tau(\mathbf{v}_E) \bmod q$. We analyse $\mathbf{u}$ reduced modulo q since is has the shortest norm in the lattice. Consider the following cases. Firstly, assume $\mathbf{v}_A \neq 0 \bmod q$. The proof of [SS11] from Lemma 26[9] still applies with the shift $\mathrm{Rot}(\mathbf{E})^T \cdot \tau(\mathbf{v}_E) \bmod q$ and only adds another union bound over $\mathbf{v}_E \in R_q^l$. Therefore, for any $\varepsilon' > 0$ in this case we have,

$$\Pr\left(\|\tau(\mathbf{u})\| \geq \frac{q}{\sqrt{N} e^{\varepsilon'} \cdot q^{(n+l)/m}} \right) \geq 1 - 2^{-Nm(\varepsilon' - 3)}.$$

Otherwise, if $\mathbf{v}_A = 0 \bmod q$ and $0 < \|\tau(\mathbf{v}_E) \bmod q\| \leq q/2B_{\max}$

$$\|\tau(\mathbf{u})\| = \left\|\mathrm{Rot}(\mathbf{E})^T \cdot \tau(\mathbf{v}_E)\right\| \leq \mathfrak{s}_1(\mathrm{Rot}(\mathbf{E})^T) \cdot \|\tau(\mathbf{v}_E)\| < B_{\max} \cdot q/2B_{\max} = q/2.$$

This means that $\mathrm{Rot}(\mathbf{E})^T \cdot \tau(\mathbf{v}_E)$ is never reduced modulo q and it belongs to the set $\mathrm{Rot}(\mathbf{E})^T \cdot \mathbb{Z}^{Nl}$ we excluded from the lattice. Lastly, if $\mathbf{v}_A = 0 \bmod q$ and $q/2B_{\max} < \|\tau(\mathbf{v}_E) \bmod q\|$, we lower bound the norm

$$\|\tau(\mathbf{u})\| = \left\| \begin{bmatrix} \mathbf{I}_{Nl} \\ \mathrm{Rot}(\bar{\mathbf{E}})^T \end{bmatrix} \cdot \tau(\mathbf{v}_E) \bmod q \right\| \geq \|\tau(\mathbf{v}_E) \bmod q\| > \frac{q}{2B_{\max}}.$$

Overall, the shortest vector is larger than $q \cdot \min\left(\frac{1}{2B_{\max}}, (\sqrt{N} e^{\varepsilon'} \cdot q^{(n+l)/m})^{-1}\right)$. Combining this bound with the arguments above we get

$$\begin{aligned} \lambda_{N(m-l)}(\tau(\Lambda)) &\leq \frac{Nm}{1/q \cdot \lambda_{Nl+1}(\Lambda_q(\mathbf{M}^T))} \\ &\leq Nm \cdot \max\left(2B_{\max}, \sqrt{N} e^{\varepsilon'} \cdot q^{(n+l)/m}\right). \end{aligned} \tag{5}$$

For vectors $\mathbf{x}$ with norm bounded this way it holds

$$\begin{aligned} \|\mathrm{Rot}(\mathbf{E}) \cdot \tau(\mathbf{x})\| &\leq \mathfrak{s}_1(\mathrm{Rot}(\mathbf{E})) \cdot \|\tau(\mathbf{x})\| \\ &< B_{\max} \cdot Nm \cdot \max\left(2B_{\max}, \sqrt{N} e^{\varepsilon'} \cdot q^{(n+l)/m}\right) \leq q/2 \end{aligned}$$

by our conditions on q. In other words, $\mathbf{Ex} = \mathbf{0}$ also holds over the integers. Therefore, $\lambda_{N(m-l)}(\tau(\Lambda_q^{\perp}(\mathbf{A}) \cap \Lambda^{\perp}(\mathbf{E}))) \leq \lambda_{N(m-l)}(\tau(\Lambda))$ from Equation (5). Finally, when $(\mathbf{A}, \mathbf{E})$ are not necessarily in $\mathcal{S}$ we adapt the statistical distance in the same way as in Theorem 6, concluding the proof. □

Remark 3. Since the proof relies on the smoothing argument and Corollaries 1+2, the transformation between embeddings from Corollary 8 still applies here. For the canonical embedding statement we require the same conditions and

$$\mathbf{\Sigma} \geq \mathfrak{s}_1(\nu) \cdot Nm \cdot \max\left(2B_{\max}, \sqrt{N} e^{\varepsilon'} \cdot q^{(n+l)/m}\right) \cdot \sqrt{\ln(2N(m-l)(1+1/\varepsilon))/\pi}.$$

Remark 4. For Condition 1 need $\mathfrak{s}_1(\mathrm{Rot}(\mathbf{E})) < B_{\max}$ so the distribution of $\bar{\mathbf{E}}$ has to have short outputs. For Condition 2, every vector $\mathbf{w} = (\mathbf{w}_1, \mathbf{w}_2)$ in $\Lambda^{\perp}(\mathbf{E})$ can be represented as $\mathbf{w}_1 = -\bar{\mathbf{E}} \cdot \mathbf{w}_2$ for some $\mathbf{w}_2 \in R^{m-l}$. Hence, the matrix product is equal to

$$[\mathbf{A}_1 \mid \mathbf{A}_2] \cdot \begin{bmatrix} -\bar{\mathbf{E}} \cdot \mathbf{w}_2 \\ \mathbf{w}_2 \end{bmatrix} = (-\bar{\mathbf{E}} \cdot \mathbf{A}_1 + \mathbf{A}_2) \cdot \mathbf{w}_2.$$

The matrix $\mathbf{A}_2 \in R_q^{n \times (m-l)}$ is uniform and independent of the other summand, so the distribution of the matrix sum is still uniform. Then by Lemma 21 we can take $1 - \delta_2 = \prod_{i=0}^{n-1}\left(1 - \frac{1}{q^{\delta(m-l-i)}}\right)^f$.

[9] Here the proof strategy only applies to the power-of-2 cyclotomics in the coefficient embedding.

4.4 Smoothing in the Presence of Noisy Linear Leakage

Lastly, we consider leakage of the form $(\mathbf{E}, \mathbf{E} \cdot \mathbf{x} + \mathbf{y})$ for a given matrix $\mathbf{E}$ and a hidden noise vector $\mathbf{y}$. The matrix $\mathbf{E}$ can be sampled from a short distribution or be fixed as long as it satisfies the conditions below. Moreover, we allow for Q many leaked hints defined by different matrices $(\mathbf{E}_i)_{i\in[Q]}$ and noise vectors $(\mathbf{y}_i)_{i\in[Q]}$ coming from potentially different noise distributions.

Our key observation is that recent advancements in studying Module Learning With Errors under noisy leakage are also useful in the context of our smoothing LHL under leakage. We recall the following result.

Lemma 27 ([ENP24, Lem. 1]). *Let m, Q be positive integers, and $\sigma_{\mathbf{x}}, (\sigma_{\mathbf{y},i})_{i\in[Q]}$ be positive reals. Take $\mathbf{E}_0, \ldots, \mathbf{E}_{Q-1} \in \mathbb{Z}^{m\times m}$ and let $\mathbf{I}$ denote the identity matrix of dimension m. We define $\mathbf{\Sigma} = \left(\frac{1}{\sigma_{\mathbf{x}}^2}\mathbf{I} + \sum_{i\in[Q]} \frac{1}{\sigma_{\mathbf{y},i}^2} \cdot \mathbf{E}_i^T\mathbf{E}_i\right)^{-1}$. Then the following two distributions over $\mathbb{Z}^{m(Q+1)}$ are identical:*

$$D_1 = \left\{(\mathbf{x}, \mathbf{z}_0, \ldots, \mathbf{z}_{Q-1}) | \mathbf{x} \leftarrow \mathcal{D}_{\mathbb{Z}^m,\sigma_{\mathbf{x}}}, \mathbf{y}_i \leftarrow \mathcal{D}_{\mathbb{Z}^m,\sigma_{\mathbf{y}_i}}, \mathbf{z}_i = \mathbf{E}_i \cdot \mathbf{x} + \mathbf{y}_i\right\}$$

$$D_2 = \left\{ (\widehat{\mathbf{x}}, \mathbf{z}_0, \ldots, \mathbf{z}_{Q-1}) \middle| \begin{array}{l} \mathbf{x} \leftarrow \mathcal{D}_{\mathbb{Z}^m,\sigma_{\mathbf{x}}}, \mathbf{y}_i \leftarrow \mathcal{D}_{\mathbb{Z}^m,\sigma_{\mathbf{y}_i}}, \mathbf{z}_i = \mathbf{E}_i \cdot \mathbf{x} + \mathbf{y}_i \\ \mathbf{c} = \mathbf{\Sigma} \cdot \sum\limits_{i\in[Q]} \dfrac{1}{\sigma_{\mathbf{y},i}^2} \mathbf{E}_i^T \mathbf{z}_i, \quad \widehat{\mathbf{x}} \leftarrow \mathcal{D}_{\mathbb{Z}^m,\sqrt{\mathbf{\Sigma}},\mathbf{c}} \end{array} \right\}$$

The above lemma allows us to derive the following smoothing-based LHL with noisy linear leakage.

Theorem 8. *Let K be a cyclotomic field of degree N. Let q be an integer. Let $\mathcal{P}_A$ be a distribution on the set of matrices $R_q^{n\times m}$ for $n \le m$. Let $\varepsilon \in (0,1)$, and $\beta, \delta_1, \delta_2 > 0$. Suppose that*

$$\Pr\left[\eta_\varepsilon(\Lambda_q^\perp(\mathbf{A})) > \beta | \mathbf{A} \leftarrow \mathcal{P}_A\right] < \delta_1,$$
$$\Pr\left[\mathrm{im}(\mathbf{A}) \neq R_q^n | \mathbf{A} \leftarrow \mathcal{P}_A\right] < \delta_2.$$

Let $Q > 0$ be an integer, and $\sigma_{\mathbf{x}}, (\sigma_{\mathbf{y},i})_{i\in[Q]}$ be positive reals. For $i \in [Q]$ take $\mathbf{E}_i \in R^{m\times m}$, defining its corresponding rotation matrix $\mathrm{Rot}(\mathbf{E}_i) \in \mathbb{Z}^{mN\times mN}$, and let $\mathbf{I}$ denote the identity matrix of dimension Nm. We define

$$\mathbf{\Sigma} = \left(\frac{1}{\sigma_{\mathbf{x}}^2}\mathbf{I} + \sum_{i\in[Q]} \frac{1}{\sigma_{\mathbf{y},i}^2} \cdot \mathrm{Rot}(\mathbf{E}_i)^T \,\mathrm{Rot}(\mathbf{E}_i)\right)^{-1},$$

and assume that $\mathbf{\Sigma} \ge \beta^2$. Define

$$\mathcal{P}' = \left\{(\mathbf{A}\mathbf{x} \bmod q, \mathbf{A}, (\mathbf{E}_i, \mathbf{E}_i\mathbf{x} + \mathbf{y}_i)_i) \mid \mathbf{A} \leftarrow \mathcal{P}_A, \mathbf{x} \leftarrow \mathcal{D}_{R^m,\sigma_{\mathbf{x}}}, \mathbf{y}_i \leftarrow \mathcal{D}_{R^m,\sigma_{\mathbf{y}_i}}\right\}$$
$$\mathcal{U}' = \left\{(\mathbf{u} \leftarrow \mathcal{U}(R_q^n), \mathbf{A}, (\mathbf{E}_i, \mathbf{E}_i\mathbf{x} + \mathbf{y}_i)_i) \mid \mathbf{A} \leftarrow \mathcal{P}_A, \mathbf{x} \leftarrow \mathcal{D}_{R^m,\sigma_{\mathbf{x}}}, \mathbf{y}_i \leftarrow \mathcal{D}_{R^m,\sigma_{\mathbf{y}_i}}\right\}$$

where $i \in [Q]$. Then

$$\mathrm{SD}(\mathcal{P}', \mathcal{U}') < \varepsilon/(1-\varepsilon) + \delta_1 + \delta_2, \qquad \mathrm{RD}(\mathcal{P}'; \mathcal{U}') < \frac{1+\varepsilon}{1-\varepsilon} + (\delta_1 + \delta_2) \cdot q^{Nn}.$$

Proof. Let D_1 and D_2 be the two distributions from Lemma 27 using the rotation matrices $\mathrm{Rot}(\mathbf{E}_0), \ldots, \mathrm{Rot}(\mathbf{E}_{Q-1})$. Let us rewrite $\mathcal{P}'$ and introduce $\mathcal{P}''$ as follows.

$$\mathcal{P}' = \left\{(\mathbf{A}\mathbf{x} \bmod q, \mathbf{A}, (\mathbf{E}_i, \mathbf{z}_i)_{i\in[Q]}) \mid \mathbf{A} \leftarrow \mathcal{P}_A, \tau(\mathbf{x}, \mathbf{z}_0, \ldots, \mathbf{z}_{Q-1}) \leftarrow D_1\right\},$$

$$\mathcal{P}'' = \left\{ (\mathbf{A}\widehat{\mathbf{x}} \bmod q, \mathbf{A}, (\mathbf{E}_i, \mathbf{z}_i)_{i\in[Q]}) \mid \mathbf{A} \leftarrow \mathcal{P}_A, \tau(\widehat{\mathbf{x}}, \mathbf{z}_0, \ldots, \mathbf{z}_{Q-1}) \leftarrow D_2 \right\}.$$

By Lemma 27, the two distributions are identical, hence not changing the SD or RD value. In $\mathcal{P}''$, the distribution of the hash input $\widehat{\mathbf{x}}$ comes from a discrete Gaussian distribution with covariance matrix $\mathbf{\Sigma}$ defined in the lemma statement and with center $\tau^{-1}(\mathbf{c})$, where $\mathbf{c} = \mathbf{\Sigma} \cdot \sum_{i\in[Q]} \frac{1}{\sigma^2_{\mathbf{y},i}} \operatorname{Rot}(\mathbf{E_i})^T \tau(\mathbf{z})_i$. As by assumption $\mathbf{\Sigma} \geq \beta^2$, we can now invoke Theorem 6, concluding the proof. □

5 Properties of Some Concrete Distributions

The different LHL flavours we presented in Section 3.3 make use of the collision probability or well-spreadness of a given hash input distribution $\mathcal{P}$ over R_q^m for a ring R, modulus q and dimension m. In some cases, we might assume that $\mathcal{P}$ is obtained by independently and identically sampling every of the m ring elements over some distribution $\mathcal{X}$ over R_q, that is, $\mathcal{P} = \mathcal{X}^m$. In this section, we give bounds on well-spreadness and collision probability for different distributions over cyclotomic rings which appeared in the lattice literature. Some bounds have been (implicitly) proven before, some are generalised results and some are new bounds. We will detail our contribution for every distribution. All results are stated with respect to the coefficient embedding.

Bounded Uniform Distribution The following lemma gives concrete formulas for the collision probability and the well-spreadness of the bounded uniform distribution over R_q. All were already implicitly proved in [Mic07]. The bounded uniform distribution is used in many places in lattice-based cryptography. A popular example is the signature scheme Dilithium [LDK+22].

Lemma 28 (Adapted from [Mic07]). *Let R be a cyclotomic ring of degree N and let q be an unramified prime number such that $\Phi(X) = \prod_{j=1}^{f} \Phi_j(X) \bmod q$ is the factorisation of the cyclotomic polynomial in R_q with every Φ_j being of degree δ and $N = f\delta$. Let $\eta < q/2$ be a positive integer. Let $\mathcal{I} = \langle \prod_{j\in G} \Phi_j(X) \rangle$ for some subset $G \subset \{1, \ldots, f\}$ be an arbitrary ideal of R_q (this respresentation of $\mathcal{I}$ is correct by Lemma 4). Let $\mathcal{X}$ be the uniform distribution over $S_\eta = \{c \in R_q \colon \|\tau(c)\|_\infty \leq \eta\}$, defining the hash input distribution $\mathcal{P} = \mathcal{X}^m$ for some dimension m. Then*

- $\operatorname{Coll}(\mathcal{X} \bmod \mathcal{I}) \leq 1/(2\eta+1)^{\delta|G|}$,
- $\operatorname{Coll}(\mathcal{P} \bmod \mathcal{I}) \leq 1/(2\eta+1)^{\delta|G|m}$,
- $\operatorname{Coll}(\mathcal{P}) = 1/(2\eta+1)^{Nm}$.

In particular the distribution $\mathcal{X}$ is $1/(2\eta+1)^\delta$-well-spread.

Proof. Firstly, $\operatorname{Coll}(\mathcal{P}) = \operatorname{Coll}(\mathcal{X})^m$, where

$$\operatorname{Coll}(\mathcal{X}) = \sum_{c\in S_\eta} \Pr[c \leftarrow \mathcal{X}]^2 = (2\eta+1)^N \cdot \frac{1}{(2\eta+1)^{2N}} = \frac{1}{(2\eta+1)^N}.$$

For the collision probability of the reduced distribution, let $c(X) \leftarrow \mathcal{X}$. Then

$$c(X) = \underbrace{\sum_{i=0}^{d-1} \alpha_i X^i}_{c_0} + \underbrace{\sum_{i=d}^{N-1} \alpha_i X^i}_{c_1},$$

where all $\alpha_i \leftarrow \mathcal{U}(\{-\eta, \ldots, \eta\})$. For any fixed value of c_1, the sum $c_0 + c_1 = c$ is distributed uniformly in a subset of R_q of size $(2\eta+1)^{\delta|G|}$ dependent on the value of $(c_1 \bmod f_I)$. Then for two samples $c_0 + c_1 \leftarrow \mathcal{X}, c_0' + c_1' \leftarrow \mathcal{X}$

$$\begin{aligned}\Pr(c_0 + (c_1 \bmod f_I) = c_0' + (c_1' \bmod f_I)) &= \Pr(c_0 = c_0' + (c_1' - c_1 \bmod f_I)) \\ &\leq (2\eta+1)^{-\delta|G|}.\end{aligned}$$

□

Remark 5. We can clearly see that for $\delta \neq N$, Corollary 6 provides tighter bounds than Corollary 7. Hence, it depends on the splitting behavior of the ring, which statement gives the better results.

Discrete Gaussian Distribution Second, we look at the maybe most natural distribution arising in lattice-based cryptography: discrete Gaussian distributions. Using the smoothing properties of Gaussians, we compute in Lemma 29 a novel well-spreadness bound. As Section 3 applies to distributions over R_q, we consider the discrete Gaussian distribution reduced modulo q.

Lemma 29. *Let K be a cyclotomic field with R its ring of integers. Further, let q be an unramified prime modulus such that $R_q = \prod_{j=1}^{f} \mathbb{Z}_q[X]/\langle \Phi_j(X)\rangle$ for distinct, irreducible $\Phi_j(X)$ of degree δ. Let $\varepsilon = 2^{-N}$, where N is the ring degree of R. Let $\sigma > N \cdot q^{\delta/N}$. Then $\mathcal{D}^{\theta}_{R,\sigma} \bmod q$ is B-well-spread with*

$$B = \frac{1}{q^\delta}\left(1 + 2^{-N+2}\right).$$

Proof. By Lemma 18 and Lemma 2, for all $j \in [f]$,

$$\eta_\varepsilon(\langle q, \Phi_j(X)\rangle) \leq (N(\langle q, \Phi_j(X)\rangle) \cdot |\Delta_K|)^{1/N} \leq q^{\delta/N} \cdot N.$$

Thus $\sigma > \max_j(\eta_\varepsilon(\langle q, \Phi_j(X)\rangle))$ and by Lemma 15 $\forall y \in \mathbb{Z}_q[X]/\langle \Phi_j(X)\rangle$

$$\max_{j \in [f]} (\Pr[x \bmod (q, \Phi_j(X)) = y | x \leftarrow \mathcal{D}_{R,\sigma}]) \leq \frac{1}{q^\delta}(1 + 2^{-N+2}).$$

□

We further state a recent result on the min-entropy of discrete Gaussian distributions modulo ideals of R_q in the coefficient embedding. We state a slightly generalised version considering discrete Gaussians over R^m for $m \geq 1$.

Lemma 30 (Adapted from [JLWG25, Cor. 4]). *Let K be a cyclotomic field with R its ring of integers. Further, let q be an unramified prime modulus such that $R_q = \prod_{j=1}^{f} \mathbb{Z}_q[X]/\langle \Phi_j(X)\rangle$ for distinct, irreducible $\Phi_j(X)$ of degree δ. Let $\mathcal{I}$ be an ideal of R_q of norm $q^{\delta k}$ for some $k \in \{0, \ldots, f\}$ (this is correct by Lemma 4). For $m \in \mathbb{N}$, let σ be a positive real and $\mathbf{c} \in R^m$ be a centre. If $\sigma \leq \frac{\sqrt{\pi}}{2} \cdot \frac{q}{\sqrt{\ln(4m\delta k)}}$ it holds*

$$\mathrm{H}_\infty(\mathcal{D}^{\tau}_{R^m,\sigma,\mathbf{c}} \bmod \mathcal{I}) \geq m\delta k \log \sigma - 1.$$

Biased Ternary Distribution Next we consider biased ternary distributions. Such distributions are for instance relevant in compact proof systems and were used in [ALS20]. Let $\mathcal{D}$ be the ternary distribution with bias $p \in [0,1]$ over $\{\pm 1, 0\}$, which samples 0 with probability p and ± 1 with probability $(1-p)/2$. The following lemma provides a formula for the collision probability as well as for well-spreadness. The proof of the latter follows the proof technique of [ALS20], and generalises their result form power-of-2 to arbitrary cyclotomics. It requires the special relation from Lemma 3 between modulus q and cyclotomic polynomial $\Phi(X)$. We include the full proof below to showcase the proof technique.

Lemma 31 (Generalised from [ALS20, Lemma 3.2]). *Let $K = \mathbb{Q}[X]/\langle \Phi(X) \rangle$ be the ν-th cyclotomic field of degree $N = \varphi(\nu)$, with R its ring of integers. Further, let $\nu = \prod p_i^{e_i}$ for $e_i \geq 1$ and let $\mu = \prod p_i^{f_i}$ for $1 \leq f_i \leq e_i$. Let q be a prime such that $q = 1 \bmod \mu$ and the multiplicative order modulo ν of q is ν/μ. By Lemma 3, $\Phi(X) = \prod_{j=1}^{f}(X^\delta - r_j) \bmod q$ for distinct $r_j \in \mathbb{Z}_q^\times$ with $N = f\delta$. Let $\mathcal{X}$ be the distribution over $S_1 = \{c \in R_q \colon \|\tau(c)\|_\infty \leq 1\}$ where each coefficient of the polynomial is sampled from $\mathcal{D}$, that is, $\mathcal{X} = \mathcal{D}^N$. This defines the hash input distribution $\mathcal{P} = \mathcal{X}^m = \mathcal{D}^{Nm}$ for some dimension m. The distribution $\mathcal{X}$ is B-well-spread with*

$$B = \max_{j \in [f]} \left(\frac{1}{q} + \frac{1}{q} \sum_{k=1}^{q-1} \prod_{l=0}^{f-1} \left| p + (1-p)\cos(2\pi k r_j^l / q) \right| \right)^\delta, \tag{6}$$

and the collision probability of $\mathcal{P}$ is given by

$$\mathrm{Coll}(\mathcal{P}) = \left(\frac{1 - 2p + 3p^2}{2} \right)^{Nm}.$$

Proof. We start by proving the collision probability formula. Note that $\mathrm{Coll}(\mathcal{P}) = \mathrm{Coll}(\mathcal{D}^{Nm}) = \mathrm{Coll}(\mathcal{D})^{Nm}$, where

$$\mathrm{Coll}(\mathcal{D}) = \sum_{j \in \{\pm 1, 0\}} \mathcal{D}(j)^2 = p^2 + 2 \cdot \left(\frac{p-1}{2} \right)^2 = \frac{3p^2 - 2p + 1}{2}.$$

We now prove the well-spreadness, which requires to analyze the distribution of the CRT-slots $\bmod\ (X^\delta - r_j)$. We start by introducing some notations. Let $c(X) = \sum_{i=0}^{N-1} c_i X^i \leftarrow \mathcal{X}$. Let $S_1^{(\delta)} \subset R_q$ be the set given as

$$S_1^{(\delta)} = \Big\{ c(X) = c_0 + c_\delta X^\delta + \cdots + c_{\delta(f-1)} X^{(f-1)\delta} \mid \|c\|_\infty \leq 1 \Big\},$$

where δ divides the degree of each monomial. Then we can write S_1 as

$$S_1 = \left\{ c(X) = \tilde{c}_0(X) + \cdots + \tilde{c}_{\delta-1}(X) \cdot X^{\delta-1} \mid \tilde{c}_i(X) = \sum_{l=0}^{f-1} c_{\delta l + i} X^{\delta l} \in S_1^{(\delta)} \right\}.$$

Computing $c(X) \bmod (X^\delta - r_j)$ is equivalent to evaluating δ-powers of X as $X^\delta = r_j$. Or evaluating polynomials $\bar{c}_i(X)$ in r_j where

$$\bar{c}_i(X) \coloneqq \tilde{c}_i(X^{1/\delta}) = c_i + c_{\delta+i} X + \cdots + c_{\delta(f-1)+i} X^{f-1}.$$

Overall we obtain

$$c(X) \bmod (X^\delta - r_j) = \bar{c}_0(r_j) + \bar{c}_1(r_j) \cdot X + \cdots + \bar{c}_{\delta-1}(r_j) \cdot X^{\delta-1}.$$

Note that the distributions of each coefficient $\bar{c}_i(r_j)$ are i.i.d. and equal to $\sum_{l=0}^{f-1} c_{\delta l+i} r_j^l$, a convolution of f independent random variables $c_{\delta l+i} r_j^l$. Since distributions of $\bar{c}_i(r_j)$ are identical we prove the statement only for $i = 0$.

Denote P_1 the distribution of $\bar{c}_0(r_j) = \sum_{l=0}^{f-1} c_{\delta l} r_j^l$ with values in $\mathbb{Z}_q$. By the Fourier transform inversion property for finite fields $\forall y \in \mathbb{Z}_q$:

$$P_1(y) = \frac{1}{q} \sum_{k=0}^{q-1} \hat{P}_1(k) \cdot e^{-2\pi i y \frac{k}{q}},$$

where by $\hat{F}$ we denote the Fourier transform of a function F. For $l \in [f]$ let μ_l denote the distribution $\mu_l = \left\{ c_{\delta l} r_j^l \mid c_{\delta l} \leftarrow \mathcal{D} \right\}$. Then $P_1 = \sum_{l=0}^{f-1} \mu_l$ and $\hat{P}_1 = \prod_{l=0}^{f-1} \hat{\mu}_l$. For any $k \in \mathbb{Z}_q$

$$\begin{aligned} \hat{\mu}_l(k) &= p + \frac{1-p}{2} \cdot e^{-2\pi i k \frac{r_j^l}{q}} + \frac{1-p}{2} \cdot e^{+2\pi i k \frac{r_j^l}{q}} \\ &= p + (1-p) \cdot \cos(2\pi k r_j^l / q), \end{aligned}$$

where the transformation uses the Euler formula and the odd and even properties of the sine and cosine functions accordingly. Then

$$\begin{aligned} P_1(y) &= \frac{1}{q} \sum_{k=0}^{q-1} \prod_{l=0}^{f-1} \left(p + (1-p) \cdot \cos(2\pi k r_j^l / q) \right) \cdot e^{-2\pi i y \frac{k}{q}} \\ &\leq \frac{1}{q} + \frac{1}{q} \sum_{k=1}^{q-1} \prod_{l=0}^{f-1} \left| p + (1-p) \cdot \cos(2\pi k r_j^l / q) \right|, \end{aligned}$$

where the first inequality follows from taking the absolute value on both sides and applying the triangle inequality and the last equality only separates the term corresponding to $k = 0$. As all coefficients are identically and independently distributed, the same bound holds for all of the δ coefficients. Given that the bound on P_1 depends on the underlying r_j, the final bound on the formula of the well-spreadness B takes the maximum over all r_j. □

Remark 6. As already observed in [ALS20, Lemma 3.3], in the special case of power-of-2 cyclotomic rings, we can exploit symmetries in order to simplify the well-spreadness formula. More precisely, for any $j \in [f]$, the element r_j generates the same group. That is, $\forall j : \left\{ \pm 1, \pm r_j, \ldots, \pm r_j^{f-1} \right\} = \langle r_j \rangle = G_\mu$ where G_μ is the only subgroup of order μ in $\mathbb{Z}_q^*$. This group contains a subgroup $\langle 1, -1 \rangle$ therefore the following set $\forall j : \left\{ 1, r_j, \ldots, r_j^{f-1} \right\} = G_\mu / \langle -1 \rangle$ also forms the same quotient group for every j. This is exactly the set that appears in the product of the well-spreadness formula. Therefore, the product contains the same values for every r_j only multiplied in a different order and we don't need to take the maximum over $j \in [f]$ to obtain the bound on well-spreadness.

Similarly, we don't have to sum over all $k \in 1 \ldots q-1$. We obtain the same product value for every $k' = \alpha \cdot k$ with $\alpha \in G_\mu$. Then for all k, α we have

$$\left\{ k \cdot \alpha, k \cdot \alpha \cdot r_j, \ldots, k \cdot \alpha \cdot r_j^{f-1} \right\} = k \cdot (\alpha \cdot G_\mu) / \langle -1 \rangle = k \cdot G_\mu / \langle -1 \rangle.$$

Then the sum becomes

$$P_1(y) = \frac{1}{q} + \frac{2f}{q} \sum_{k \in \mathbb{Z}_q^* / G_\mu} \prod_{l=0}^{f-1} \left| p + (1-p) \cdot \cos(2\pi k r_j^l / q) \right|,$$

and numerical estimations can be computed $2f$ times faster. The same optimisations apply to Lemma 32 and Lemma 33 in the case of power-of-2 cyclotomics. The group arguments above are not true for arbitrary cyclotomic rings, as for example -1 may not even belong to G_μ.

Fixed Hamming Weight Ternary Distribution Another distribution that appears in the literature, is the uniform bounded distribution with an additional restriction on the Hamming weight of the vectors. Similar to the biased ternary distribution, it appears in the context of compact proof systems. It allows for even more efficient proofs as it generates elements of smaller ℓ_1-norm. Concretely, we consider the uniform distribution over

$$S_{1,w} = \left\{ c(X) = \tilde{c}_0(X) + \cdots + \tilde{c}_{\delta-1}(X) \cdot X^{\delta-1} \mid \tilde{c}_i \in S_1^{(\delta)} \wedge \mathrm{HW}(\tilde{c}_i) = \tilde{w} \right\}.$$

Note that every element $c(X)$ of $S_{1,w}$ has the same Hamming weight $\mathrm{HW}(c) = \delta \cdot \tilde{w} := w$ and coefficients equal to ± 1 or 0. We remark that the non-zero coefficients are somewhat regularly distributed. This tweak simplifies the analysis of well-spreadness significantly.

The same proof technique applies when the non-zero coefficients are distributed in the set $S_\eta \setminus \{0\}$, although it makes the Fourier Transform of the coefficient distributions more complex. A similar distribution was analysed in [ESLR23].

Below, we provide a formula for the collision probability and the well-spreadness. The latter has been proven for power-of-2 cyclotomics in the work of [ESZ22], where they built upon the approach of [ALS20] detailed above. The main obstacle the proof deals with is that the Hamming weight restriction introduces dependencies into some of the distributions. The authors first remove this dependency and apply the Fourier analysis tools and then prove that the distribution analysed is close enough to the original one. We generalise the result to arbitrary cyclotomics and provide a proof for completeness. As for the biased ternary case, we require the special relation from Lemma 3 between modulus q and cyclotomic polynomial $\Phi(X)$. The full proof is provided in Appendix A.1

Lemma 32 (Generalised from [ESZ22, Lemma 1]). *Let $K = \mathbb{Q}[X]/\langle\Phi(X)\rangle$ be the ν-th cyclotomic field of degree $N = \varphi(\nu)$, with R its ring of integers. Further, let $\nu = \prod p_i^{e_i}$ for $e_i \geq 1$ and let $\mu = \prod p_i^{f_i}$ for $1 \leq f_i \leq e_i$. Let q be a prime such that $q = 1 \bmod \mu$ and the multiplicative order modulo ν of q is ν/μ. By Lemma 3, $\Phi(X) = \prod_{j=1}^{f}(X^\delta - r_j) \bmod q$ for distinct $r_j \in \mathbb{Z}_q^\times$ with $N = f\delta$. Let $\mathcal{X}$ be the uniform distribution over $S_{1,w}$, defining the hash input distribution $\mathcal{P} = \mathcal{X}^m$ for some dimension m. The distribution $\mathcal{X}$ is B-well-spread with*

$$B = \max_{j\in[f]} \left(\frac{\beta}{q} + \frac{\beta}{q} \sum_{k=1}^{q-1} \left| \frac{1}{f} \sum_{l=0}^{f-1} \cos(2\pi k r_j^l / q) \right|^{\tilde{w}} \right)^\delta, \tag{7}$$

where $\beta = \frac{f^{\tilde{w}}(f-\tilde{w})!}{f!}$ and the collision probability of $\mathcal{P}$ is given by

$$\mathrm{Coll}(\mathcal{P}) = \left(\binom{f}{\tilde{w}}^{-\delta} 2^{-w} \right)^m.$$

Centered Binomial Uniform Distribution Let $\mathcal{D}$ be the centered binomial distribution with parameter η, defined over $j \in \{-\eta, \ldots, \eta\}$ by

$$\Pr[j \leftarrow \mathcal{D}] = 2^{-2\eta} \cdot \binom{2\eta}{\eta + j}.$$

By sampling each coefficient independently over $\mathcal{D}$, we obtain a distribution $\mathcal{X}$ over the ring R of degree N, where $\mathcal{X} = \mathcal{D}^N$. Similar to the bounded uniform distribution, it has found widespread use in lattice-based cryptography. For instance, it is used in the encryption scheme Kyber [SAB+22]. We provide a formula of collision probability and well-spreadness in the lemma below, applying to arbitrary cyclotomic rings. The latter has been analysed in [CLS16] for the special case of power-of-2 cyclotomics. The detailed proof can be found in Appendix A.2.

Lemma 33. *Let $K = \mathbb{Q}[X]/\langle \Phi(X) \rangle$ be the ν-th cyclotomic field of degree $N = \varphi(\nu)$, with R its ring of integers. Further, let $\nu = \prod p_i^{e_i}$ for $e_i \geq 1$ and let $\mu = \prod p_i^{f_i}$ for $1 \leq f_i \leq e_i$. Let q be a prime such that $q = 1 \bmod \mu$ and the multiplicative order modulo ν of q is ν/μ. By Lemma 3, $\Phi(X) = \prod_{j=1}^{f} (X^\delta - r_j) \bmod q$ for distinct $r_j \in \mathbb{Z}_q^\times$ with $N = f\delta$. Define $\mathcal{X} = \left\{ c(X) = \sum_{i=0}^{N-1} c_i X^i \mid c_i \leftarrow \mathcal{D} \right\}$, that is, $\mathcal{X} = \mathcal{D}^N$. This defines the hash input distribution $\mathcal{P} = \mathcal{X}^m = \mathcal{D}^{Nm}$ for some dimension m.*

The collision probability of $\mathcal{P}$ is given by

$$\mathrm{Coll}(\mathcal{P}) = \left(2^{-4\eta} \binom{4\eta}{2\eta} \right)^{Nm},$$

and the distribution $\mathcal{X}$ is B-well spread with

$$B = \max_{j \in [f]} \left(\frac{1}{q} + \frac{1}{q} \sum_{k=1}^{q-1} \prod_{l=0}^{f-1} |\cos(\pi k r_j^l / q)|^{2\eta} \right)^{\delta}. \tag{8}$$

Remark 7. When $\eta = 1$ we recover the result in Lemma 31 for $p = 1/2$, using $\cos(x)^2 = \frac{1}{2} + \frac{1}{2}\cos(2x)$.

References

AAB$^+$24. Marius A. Aardal, Diego F. Aranha, Katharina Boudgoust, Sebastian Kolby, and Akira Takahashi. Aggregating falcon signatures with LaBRADOR. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part I*, volume 14920 of *LNCS*, pages 71–106. Springer, Cham, August 2024. 1.1, 2.2

ACPS09. Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In Shai Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 595–618. Springer, Berlin, Heidelberg, August 2009. 1

ACX19. Thomas Attema, Ronald Cramer, and Chaoping Xing. A note on short invertible ring elements and applications to cyclotomic and trinomials number fields. Cryptology ePrint Archive, Report 2019/1200, 2019. 1

AGHS13. Shweta Agrawal, Craig Gentry, Shai Halevi, and Amit Sahai. Discrete Gaussian leftover hash lemma over infinite domains. In Kazue Sako and Palash Sarkar, editors, *ASIACRYPT 2013, Part I*, volume 8269 of *LNCS*, pages 97–116. Springer, Berlin, Heidelberg, December 2013. 1

AKPW13. Joël Alwen, Stephan Krenn, Krzysztof Pietrzak, and Daniel Wichs. Learning with rounding, revisited - new reduction, properties and applications. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 57–74. Springer, Berlin, Heidelberg, August 2013. 1

ALS20. Thomas Attema, Vadim Lyubashevsky, and Gregor Seiler. Practical product proofs for lattice commitments. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 470–499. Springer, Cham, August 2020. 5, 31, 6, 5

BC22. Iván Blanco-Chacón. On the RLWE/PLWE equivalence for cyclotomic number fields. *Applicable Algebra in Engineering, Communication and Computing*, 33:53–71, 2022. 2.1

BD20. Zvika Brakerski and Nico Döttling. Hardness of LWE on general entropic distributions. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 551–575. Springer, Cham, May 2020. 1, 1

BdPMW16. Florian Bourse, Rafaël del Pino, Michele Minelli, and Hoeteck Wee. FHE circuit privacy almost for free. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 62–89. Springer, Berlin, Heidelberg, August 2016. 1

BI22. Florian Bourse and Malika Izabachène. Plug-and-play sanitization for TFHE. Cryptology ePrint Archive, Report 2022/1438, 2022. 1, 1.1, 3.3

BJRW23. Katharina Boudgoust, Corentin Jeudy, Adeline Roux-Langlois, and Weiqiang Wen. On the hardness of module learning with errors with short distributions. *Journal of Cryptology*, 36(1):1, January 2023. 1, 1.1, 2.2, 21, 3.3

BLL$^+$15. Shi Bai, Adeline Langlois, Tancrède Lepoint, Damien Stehlé, and Ron Steinfeld. Improved security proofs in lattice-based cryptography: Using the Rényi divergence rather than the statistical distance. In Tetsu Iwata and Jung Hee Cheon, editors, *ASIACRYPT 2015, Part I*, volume 9452 of *LNCS*, pages 3–24. Springer, Berlin, Heidelberg, November / December 2015. 2

BLP$^+$13. Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 575–584. ACM Press, June 2013. 1

BLR$^+$18. Shi Bai, Tancrède Lepoint, Adeline Roux-Langlois, Amin Sakzad, Damien Stehlé, and Ron Steinfeld. Improved security proofs in lattice-based cryptography: Using the Rényi divergence rather than the statistical distance. *Journal of Cryptology*, 31(2):610–640, April 2018. 1, 4

CLS16. Hao Chen, Kristin E. Lauter, and Katherine E. Stange. Security considerations for Galois non-dual RLWE families. In Roberto Avanzi and Howard M. Heys, editors, *SAC 2016*, volume 10532 of *LNCS*, pages 443–462. Springer, Cham, August 2016. 5

Con. Keith Conrad. The different ideal. `https://kconrad.math.uconn.edu/blurbs/gradnumthy/idealfactor.pdf`, last accessed on 17.07.2024. 3.1

CPS$^+$20. Chitchanok Chuengsatiansup, Thomas Prest, Damien Stehlé, Alexandre Wallet, and Keita Xagawa. ModFalcon: Compact signatures based on module-NTRU lattices. In Hung-Min Sun, Shiuh-Pyng Shieh, Guofei Gu, and Giuseppe Ateniese, editors, *ASIACCS 20*, pages 853–866. ACM Press, October 2020. 1

DGKS21. Dana Dachman-Soled, Huijing Gong, Mukul Kulkarni, and Aria Shahverdi. Towards a ring analogue of the leftover hash lemma. *J. Math. Cryptol.*, 15(1):87–110, 2021. 1, 1.1, 4.3

DORS08. Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM journal on computing*, 38(1):97–139, 2008. 1, 2.2, 7, 9, 3.3

ENP24. Thomas Espitau, Guilhem Niot, and Thomas Prest. Flood and submerse: Distributed key generation and robust threshold signature from lattices. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part VII*, volume 14926 of *LNCS*, pages 425–458. Springer, Cham, August 2024. 1.1, 4.3, 27

ESLR23. Muhammed F. Esgin, Ron Steinfeld, Dongxi Liu, and Sushmita Ruj. Efficient hybrid exact/relaxed lattice proofs and applications to rounding and VRFs. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part V*, volume 14085 of *LNCS*, pages 484–517. Springer, Cham, August 2023. 5

ESZ22. Muhammed F. Esgin, Ron Steinfeld, and Raymond K. Zhao. MatRiCT$^+$: More efficient post-quantum private blockchain payments. In *2022 IEEE Symposium on Security and Privacy*, pages 1281–1298. IEEE Computer Society Press, May 2022. 5, 32

FB14. Serge Fehr and Stefan Berens. On the conditional rényi entropy. *IEEE Trans. Inf. Theory*, 60(11):6801–6810, 2014. 2.2

GJK24. Phillip Gajland, Jonas Janneck, and Eike Kiltz. A closer look at falcon. Cryptology ePrint Archive, Report 2024/1769, 2024. 2.3, 4

GKPV10. Shafi Goldwasser, Yael Tauman Kalai, Chris Peikert, and Vinod Vaikuntanathan. Robustness of the learning with errors assumption. In *ICS*, pages 230–240. Tsinghua University Press, 2010. 1

GMPW20. Nicholas Genise, Daniele Micciancio, Chris Peikert, and Michael Walter. Improved discrete gaussian and subgaussian analysis for lattice cryptography. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part I*, volume 12110 of *LNCS*, pages 623–651. Springer, Cham, May 2020. 1, 14

GPV08. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 197–206. ACM Press, May 2008. 1, 1.1, 2.3

HILL99. Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999. 1, 1, 1, 3.3

ILL89. Russell Impagliazzo, Leonid A. Levin, and Michael Luby. Pseudo-random generation from one-way functions (extended abstracts). In *21st ACM STOC*, pages 12–24. ACM Press, May 1989. 1, 1, 1, 3.3

Jeu24. Corentin Jeudy. *Design of advanced post-quantum signature schemes. (Conception d'algorithmes de signatures avancées post-quantiques)*. PhD thesis, University of Rennes 1, France, 2024. https://cjeudy.github.io/assets/pub/manuscript.pdf. 2.3

JLWG25. Haoxiang Jin, Feng-Hao Liu, Zhedong Wang, and Dawu Gu. Discrete gaussians modulo sub-lattices: New leftover hash lemmas for discrete gaussians. In *PKC 2025, Part II*, LNCS, pages 301–330. Springer, Cham, May 2025. 1, 1.1, 1.1, 3.3, 3.3, 30

KY16. Shuichi Katsumata and Shota Yamada. Partitioning via non-linear polynomial functions: More compact IBEs from ideal lattices and bilinear maps. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 682–712. Springer, Berlin, Heidelberg, December 2016. 1, 1.1, 3.3

LDK$^+$22. Vadim Lyubashevsky, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Peter Schwabe, Gregor Seiler, Damien Stehlé, and Shi Bai. CRYSTALS-DILITHIUM. Technical report, National Institute of Standards and Technology, 2022. available at https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022. 5

LPR10. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 1–23. Springer, Berlin, Heidelberg, May / June 2010. 1

LPR13. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for ring-LWE cryptography. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 35–54. Springer, Berlin, Heidelberg, May 2013. 1, 1.1, 1.1, 2, 5, 6, 2.1, 13, 17, 19, 4.2, 4.2, 24, 25

LS15. Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *DCC*, 75(3):565–599, 2015. 1, 1

LS18. Vadim Lyubashevsky and Gregor Seiler. Short, invertible elements in partially splitting cyclotomic rings and applications to lattice-based zero-knowledge proofs. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 204–224. Springer, Cham, April / May 2018. 1.1, 1.1, 3

LSS14. Adeline Langlois, Damien Stehlé, and Ron Steinfeld. GGHLite: More efficient multilinear maps from ideal lattices. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 239–256. Springer, Berlin, Heidelberg, May 2014. 1

LW20. Feng-Hao Liu and Zhedong Wang. Rounding in the rings. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part II*, volume 12171 of *LNCS*, pages 296–326. Springer, Cham, August 2020. 1, 1.1, 3.1, 3.3

LWZW20. Hao Lin, Mingqiang Wang, Jincheng Zhuang, and Yang Wang. Hardness of module-LWE and ring-LWE on general entropic distributions. Cryptology ePrint Archive, Report 2020/1238, 2020. Version 2020-10-09 https://eprint.iacr.org/archive/2020/1238/20201009:113322. 1.1

LWZW24. Hao Lin, Mingqiang Wang, Jincheng Zhuang, and Yang Wang. Hardness of entropic module-lwe. *Theoretical Computer Science*, page 114553, 2024. 1, 1.1, 3.3

Mic07. Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Comput. Complex.*, 16(4):365–411, 2007. 1, 1.1, 2.2, 3.1, 3.3, 5, 28

MR07. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007. 1, 1.1, 2.3, 12, 2.3

Pei08. Chris Peikert. Limits on the hardness of lattice problems in lp norms. *computational complexity*, 17(2):300–351, 2008. 11, 16

Pei10. Chris Peikert. An efficient and parallel Gaussian sampler for lattices. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 80–97. Springer, Berlin, Heidelberg, August 2010. 2.3

PR07. Chris Peikert and Alon Rosen. Lattices that admit logarithmic worst-case to average-case connection factors. In David S. Johnson and Uriel Feige, editors, *39th ACM STOC*, pages 478–487. ACM Press, June 2007. 18

Pre17. Thomas Prest. Sharper bounds in lattice-based cryptography using the Rényi divergence. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 347–374. Springer, Cham, December 2017. 2.3

Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005. 1

RSW18. Miruna Rosca, Damien Stehlé, and Alexandre Wallet. On the ring-LWE and polynomial-LWE problems. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 146–173. Springer, Cham, April / May 2018. 1, 1.1, 2.1, 20, 4.2, 4.2, 23, 4.2, 26

SAB$^+$22. Peter Schwabe, Roberto Avanzi, Joppe Bos, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, John M. Schanck, Gregor Seiler, Damien Stehlé, and Jintai Ding. CRYSTALS-KYBER. Technical report, National Institute of Standards and Technology, 2022. available at `https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022`. 1, 1.1, 5

SS11. Damien Stehlé and Ron Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 27–47. Springer, Berlin, Heidelberg, May 2011. 1, 1.1, 4.2, 22, 4.2, 26, 4.3

SSE$^+$24. Ron Steinfeld, Amin Sakzad, Muhammed F. Esgin, Veronika Kuchta, Mert Yassi, and Raymond K. Zhao. LUNA: Quasi-optimally succinct designated-verifier zero-knowledge arguments from lattices. In Bo Luo, Xiaojing Liao, Jun Xu, Engin Kirda, and David Lie, editors, *ACM CCS 2024*, pages 3167–3181. ACM Press, October 2024. 1, 1.1, 4.3

SSTX09. Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 617–635. Springer, Berlin, Heidelberg, December 2009. 1, 1.1, 3.3

vEH14. Tim van Erven and Peter Harremoës. Rényi divergence and kullback-leibler divergence. *IEEE Trans. Inf. Theory*, 60(7):3797–3820, 2014. 2.2

A Missing Proofs from Section 5

A.1 Proof of Lemma 32

Proof. We start with the collision probability of $\mathcal{P}$, where $\mathrm{Coll}(\mathcal{P}) = \mathrm{Coll}(\mathcal{X})^m$. We compute $|S_{1,\tilde{w}}| = \binom{f}{\tilde{w}}^{\delta} 2^w$. Then, the collision probability of $\mathcal{X}$, which is the uniform distribution over $S_{1,\tilde{w}}$, is $\binom{f}{\tilde{w}}^{-\delta} 2^{-w}$.

It remains to show the well-spreadness. To obtain a uniformly random element in $S_{1,\tilde{w}}$ each polynomial $\tilde{c}_i(X)$ can be sampled independently in the following way. First, we choose the non-zero monomials $(t_1, \ldots, t_{\tilde{w}})$ as an ordered $\tilde{w}$-tuple of distinct elements from $[f]$ uniformly at random. This ensures the Hamming weight condition. Then we choose the signs of the non-zero monomials $c_{\delta t_\alpha + i} \leftarrow \mathcal{U}(\{1, -1\})$ for $\alpha = 1, \ldots, \tilde{w}$. Finally, we set $\tilde{c}_i(X) = \sum_{\alpha=1}^{\tilde{w}} c_{\delta t_\alpha + i} \cdot X^{\delta \cdot t_\alpha}$, defining $c(X) = \sum_{i=0}^{\delta-1} \tilde{c}_i(X) \cdot X^i$. Note that this method samples each possible $\tilde{c}_i(X)$ polynomial $\tilde{w}!$ times, but this does not skew the distribution.

Now reduce the polynomial $c(X)$ modulo $X^\delta - r_j$ for some $j \in [f]$. Similar to the proof of Lemma 31, the distributions of each coefficient of the resulting polynomial are independent and equal to $\bar{c}_0(r_j) = \sum_{\alpha=1}^{\tilde{w}} c_{\delta t_\alpha} \cdot r_j^{t_\alpha} \in \mathbb{Z}_q$. Denote this distribution P_2. We aim to prove that $\max_{y \in \mathbb{Z}_q}(P_2(y))$ is bounded.

Define a similar distribution P_1 where instead of choosing a tuple of distinct powers of X they are allowed to repeat. In this distribution, we sample $t_\alpha \leftarrow \mathcal{U}([f]), c_{\delta t_\alpha} \leftarrow \mathcal{U}(\{1, -1\})$ for each $\alpha = 1, \ldots, \tilde{w}$. Then, P_1 is the distribution of $\bar{c}_0(r_j) = \sum_{\alpha=1}^{\tilde{w}} c_{\delta t_\alpha} \cdot r_j^{t_\alpha} \in \mathbb{Z}_q$. This simplification allows us to represent the distribution as a sum of $\tilde{w}$ independent random variables and easily compute corresponding Fourier transforms. Using the inversion property of Fourier transforms over finite fields $\forall y \in \mathbb{Z}_q$

$$P_1(y) = \frac{1}{q} + \frac{1}{q} \sum_{k=1}^{q-1} \hat{P}_1(k) \cdot e^{-2\pi i\, y \frac{k}{q}}.$$

The distribution P_1 is a convolution of $\tilde{w}$ distributions μ_α where

$$\mu_\alpha = \left\{ c_{\delta t_\alpha} \cdot r_j^{t_\alpha} \mid c_{\delta t_\alpha} \leftarrow \{1, -1\}, t_\alpha \leftarrow [f] \right\}.$$

Note that the distribution is the same for all α. Then $\hat{P}_1(k) = \prod_{\alpha=1}^{\tilde{w}} \hat{\mu}_\alpha(k) = \hat{\mu}_1(k)^{\tilde{w}}$. Define a set $S = \left\{ \pm 1, \pm r_j, \ldots, \pm r_j^{f-1} \right\}$ then $\mu_1 = \mathcal{U}(S)$ and for any $k \in \mathbb{Z}_q$

$$\begin{aligned} \hat{\mu}_1(k) &= \sum_{x \in S} \mu_1(x) \cdot e^{-2\pi i k \frac{x}{q}} \qquad (9) \\ &= \sum_{l=0}^{f-1} \frac{1}{2f} \cdot (e^{-2\pi i k \frac{r_j^l}{q}} + e^{+2\pi i k \frac{r_j^l}{q}}) = \frac{1}{2f} \sum_{l=0}^{f-1} 2 \cos(2\pi k r_j^l / q). \end{aligned}$$

In the above, the first equality uses the shape of S and $\mu_1 = \mathcal{U}(S)$, the second equality uses the Euler formula and that cosine and sine are even and odd functions accordingly. Substituting this expression into the main formula we get

$$\begin{aligned} P_1(y) &= \frac{1}{q} + \frac{1}{q} \sum_{k=1}^{q-1} \left(\frac{1}{f} \sum_{l=0}^{f-1} \cos(2\pi k r_j^l / q) \right)^{\tilde{w}} \cdot e^{-2\pi i\, y \frac{k}{q}} \\ &\leq \frac{1}{q} + \frac{1}{q} \sum_{k=1}^{q-1} \left| \frac{1}{f} \sum_{l=0}^{f-1} \cos(2\pi k r_j^l / q) \right|^{\tilde{w}}. \end{aligned}$$

For the inequality we take the absolute value on both sides and apply the triangle inequality on the right. The left side does not change since $P_1(y)$ is a positive real number. To obtain a statement about $P_2(y)$ we compute an upper bound on $\mathrm{RD}_\infty(P_2; P_1) \coloneqq \max_{y\in\mathbb{Z}_q} \frac{P_2(y)}{P_1(y)}$. Note that we take the Rényi divergence of the infinite order as for uniform distributions this is optimal.

Denote as $\mathcal{T}$ the set of ordered $\tilde{w}$-tuples with non-repeating coordinates in $[f]$. Then $|\mathcal{T}| = \frac{f!}{(f-\tilde{w})!}$. The distributions P_1, P_2 can be seen as computing the same function on the following distributions of ordered tuples

$$\begin{aligned}
D_1 &= \{((t_1, \dots, t_{\tilde{w}}), (c_{\delta t_1}, \dots, c_{\delta t_{\tilde{w}}})) \mid t_\alpha \leftarrow \mathcal{U}([f]), c_{\delta t_\alpha} \leftarrow \mathcal{U}(\{1, -1\})\}, \\
D_2 &= \{((t_1, \dots, t_{\tilde{w}}), (c_{\delta t_1}, \dots, c_{\delta t_{\tilde{w}}})) \mid (t_1, \dots, t_{\tilde{w}}) \leftarrow \mathcal{U}(\mathcal{T}), c_{\delta t_\alpha} \leftarrow \mathcal{U}(\{1, -1\})\}.
\end{aligned}$$

Note that $\mathcal{T} \subset [f]^{\tilde{w}}$ so $\mathrm{Supp}(D_2) \subset \mathrm{Supp}(D_1)$ and divergence $\mathrm{RD}_\infty(D_2; D_1)$ is well defined. Then by the data processing inequality for Rényi divergence

$$\begin{aligned}
\mathrm{RD}_\infty(P_2; P_1) &\leq \mathrm{RD}_\infty(D_2; D_1) \\
&= \frac{|\mathrm{Supp}(D_1)|}{|\mathrm{Supp}(D_2)|} = \frac{2^{\tilde{w}} f^{\tilde{w}} (f-\tilde{w})!}{2^{\tilde{w}} f!} = \frac{f^{\tilde{w}}(f-\tilde{w})!}{f!} = \beta,
\end{aligned}$$

where the first equality holds since D_1 and D_2 are uniform distributions on their support. Therefore, by the probability preservation property of RD_∞, $\forall y : P_2(y) \leq \beta \cdot P_1(y)$ which implies the main statement. □

A.2 Proof of Lemma 33

Proof. By the definition of collision probability and by the Vandermonde identity

$$\mathrm{Coll}(\mathcal{P}) = 2^{-4\eta} \sum_{i=-\eta}^{\eta} \binom{2\eta}{\eta+i}^2 = 2^{-4\eta}\binom{4\eta}{2\eta}.$$

Similarly to the proof of Lemma 31 denote P_1 the distribution of $\bar{c}_0(r_j) = \sum_{l=0}^{f-1} c_{\delta l} r_j^l$ with values in $\mathbb{Z}_q$. By the Fourier transform inversion property for finite fields $\forall y \in \mathbb{Z}_q$:

$$P_1(y) = \frac{1}{q}\sum_{k=0}^{q-1} \hat{P}_1(k) \cdot e^{-2\pi i y \frac{k}{q}}.$$

For $l \in [f]$ denote as μ_l the distribution $\mu_l = \left\{ c_{\delta l} r_j^l \mid c_{\delta l} \leftarrow \mathcal{D} \right\}$. Then $P_1 = \sum_{l=0}^{f-1} \mu_l$ and $\hat{P}_1 = \prod_{l=0}^{f-1} \hat{\mu}_l$. For $k \in \mathbb{Z}_q$

$$\begin{aligned}
\hat{\mu}_l(k) &= \sum_{s=-\eta}^{\eta} 2^{-2\eta} \cdot \binom{2\eta}{\eta+s} e^{-2\pi i k \frac{s r_j^l}{q}} = 2^{-2\eta} \cdot e^{2\pi i k \frac{\eta r_j^l}{q}} \cdot \sum_{s=0}^{2\eta} \binom{2\eta}{s} e^{-2\pi i k \frac{s r_j^l}{q}} \\
&= 2^{-2\eta} \cdot e^{2\pi i k \frac{\eta r_j^l}{q}} \cdot \left(1 + e^{-2\pi i k \frac{r_j^l}{q}}\right)^{2\eta} \\
&= 2^{-2\eta} \cdot \left(e^{\pi i k \frac{r_j^l}{q}}\right)^{2\eta} \cdot \left(1 + e^{-2\pi i k \frac{r_j^l}{q}}\right)^{2\eta} \\
&= \left(\frac{e^{\pi i k \frac{r_j^l}{q}} + e^{-\pi i k \frac{r_j^l}{q}}}{2}\right)^{2\eta} = \cos(\pi k r_j^l / q)^{2\eta},
\end{aligned}$$

where the transformations above use the binomial and Euler formulas and the odd and even property of the sine and cosine functions accordingly. Then

$$\begin{aligned}P_1(y) &= \frac{1}{q}\sum_{k=0}^{q-1}\prod_{l=0}^{f-1}\cos(\pi k r_j^l/q)^{2\eta}\cdot e^{-2\pi i y\frac{k}{q}}\\ &\leq \frac{1}{q}\sum_{k=0}^{q-1}\prod_{l=0}^{f-1}\left|\cos(\pi k r_j^l/q)\right|^{2\eta} = \frac{1}{q}+\frac{1}{q}\sum_{k=1}^{q-1}\prod_{l=0}^{f-1}\left|\cos(\pi k r_j^l/q)\right|^{2\eta},\end{aligned}$$

where the first inequality comes from taking the absolute value on both sides and applying the triangle inequality, the second equality separates the term $k = 0$. □