

A Closer Look at Falcon

Pierre-Alain Fouque¹ , Phillip Gajland², Hubert de Groote³, Jonas Janneck⁴ , and Eike Kiltz⁴ 

¹ Université de Rennes & Institut Universitaire de France

² IBM Research – Zurich

³ ENS Paris-Saclay

⁴ Ruhr University Bochum

2nd March 2026

Abstract FALCON is a winner of NIST’s six-year post-quantum cryptography standardisation competition. Based on the celebrated full-domain-hash framework of Gentry, Peikert and Vaikuntanathan (GPV) (STOC’08), FALCON leverages NTRU lattices to achieve the most compact signatures among standardised lattice-based schemes.

Its security hinges on a Rényi divergence-based argument for Gaussian samplers. However, the GPV proof, which uses statistical distance to argue closeness of distributions, fails when applied naively to FALCON due to parameter choices resulting in statistical distances as large as 2^{-34} . Additional implementation-driven deviations from the GPV framework further invalidate the original proof, leaving FALCON without a security proof despite its selection for standardisation.

In this work, we provide the first formal security proof of FALCON in the random oracle model, achieved through a few conservative modifications, now incorporated into the forthcoming standard. At the heart of our analysis lies an adaptation of the GPV framework to work with the Rényi divergence, along with an optimised method for parameter selection under this measure. We also analyse the FFO SAMPLER that is used in FALCON. Further, we prove the equivalence of *plain unforgeability* to a multi-target inhomogeneous SIS problem, and *strong unforgeability* to a second-preimage version of this problem, providing clear targets for cryptanalysis. Assuming these problems are as hard as standard SIS, we demonstrate that FALCON-512 barely satisfies the claimed 120-bit security target, while FALCON-1024 achieves the claimed security level.

Contents

1	Introduction	3
1.1	Contributions	3
1.2	Technical Overview	5
2	Preliminaries	8
2.1	Notation	8
2.2	Signatures	9
2.3	Lattices	9
2.4	Rényi Divergence	12
2.5	Hardness Assumptions	13
3	Security arguments using the Rényi Divergence	14
4	COREFALCON ⁺ : A Framework for Falcon	16
4.1	FALCON Parameter Sets	17
4.2	Security Bounds for COREFALCON ⁺	17
5	Parameters and Analysing the Security Bound	18
5.1	Security of t - $\mathcal{R}$ -ISIS and t - $\mathcal{R}$ -SPISIS	19
5.2	Number Of Signing Repetitions C_s	19
5.3	Rényi Terms	19
5.4	Final Security and Discussion	20
A	Additional Preliminaries	26
B	Proofs for Section 2 and Section 3	26
B.1	Proof of Lemma 8	26
B.2	Proof of Lemma 9	27
B.3	Proof of Corollary 2	27
B.4	Proof of Lemma 10	28
B.5	Proof of Lemma 11	28
C	Proof of Theorem 1	28
D	Proof of Theorem 2	32
E	Appendix for Section 5	34
E.1	Proof of Lemma 12	35
E.2	Proof of Corollary 3	36
E.3	Proof of Corollary 4	37
E.4	Additional Rényi Corollaries	38
F	Samplers	38
F.1	Orthogonalizations	38
F.2	KLEIN SAMPLER	39
F.3	FFO SAMPLER	41

1 Introduction

Among the 69 submissions to the NIST post-quantum cryptography standardisation process in 2016 [Kim16], FALCON [PFH⁺20] was selected as one of four winning algorithms in 2022. Currently, NIST is in the process of drafting the corresponding FIPS standard. FALCON is a signature scheme based on the full-domain-hash (FDH) paradigm [BR96], commonly known as “*hash-and-sign*”. In this framework, the public verification key is a trapdoor permutation f and the signing key is the inverse f^{-1} . To sign a message m , one first hashes m to some point $y = H(m)$ in the range of f , then outputs the signature $\sigma = f^{-1}(y)$. Verification consists of checking that $f(\sigma) = H(m)$. FALCON, like most of the selected algorithms such as KYBER [SAB⁺22] and DILITHIUM [LDK⁺22], relies on the hardness of lattice problems. Its design follows the FDH framework over lattices, as formalised in the celebrated work of Gentry, Peikert and Vaikuntanathan (GPV) [GPV08], which generalised the FDH paradigm to work with *preimage sampleable trapdoor functions*, rather than solely permutations. Concretely, GPV signatures σ are sampled from $f^{-1}(H(m))$. By leveraging NTRU lattices, introduced by Hoffstein, Pipher, and Silverman [HPS98, HHP⁺03], FALCON benefits from their ring structure. This allows a reduction in public keys by a factor of $\mathcal{O}(n)$ and accelerates many computations by a factor of $\mathcal{O}(n/\log n)$, where n is the polynomial ring dimension. More importantly, [DLP14] showed that, by choosing appropriate parameters, the length of NTRU trapdoors can be within a small constant factor of the theoretical optimal, achieving the most compact signatures among lattice-based schemes. These optimal parameters can be efficiently generated using a key generation algorithm from [PP19], which leverages the tower-of-fields structure in powers of 2 cyclotomic fields. The final component of FALCON is an efficient sampler derived from the Fast Fourier Orthogonalization (FFO) technique described by Ducas and Prest in [DP16] that samples in time $\mathcal{O}(n \log n)$, again leveraging the tower-of-fields structure. Compared to other signature schemes selected for standardisation by NIST, such as DILITHIUM [LDK⁺22] and SPHINCS+ [HBD⁺22], FALCON stands out for its compactness, minimising both public key and signature sizes.

While the GPV framework was originally proven [GPV08] under the plain (unstructured) Short Integer Solution (SIS) assumption [Ajt96], adapting it to the (structured) NTRU-SIS setting is described in the FALCON specification as “*straightforward*”. The GPV proof relies on the “*leftover hash lemma*” [HILL99, Lem. 4.8] to argue that the simulation of the random oracle is statistically close to uniform. While this statistical argument can be adapted using a regularity lemma for rings [SS11, LPR13, RSW18], applying this argument with FALCON parameters leads to statistical distances as large as 2^{-34} . Moreover, FALCON deviates from the GPV framework by relying on the Rényi divergence instead of statistical distance, to achieve tighter parameters and smaller signature sizes. Therefore, as stated in [LAZ19, Sec. 2.3], the parameters used in FALCON are not supported by the GPV proof.

Given the importance of thoroughly understanding schemes intended for mass deployment, and in light of recent classical attacks on post-quantum schemes [Beu22, CD23, MMP⁺23, Rob23], careful security analysis is paramount. Despite successfully progressing through all three stages of the NIST process and being selected for standardisation, a formal proof of FALCON remains elusive raising the following pertinent question.

Can Falcon be proven secure? If so, what is its concrete security?

1.1 Contributions

This work provides the first concrete security analysis of FALCON-type signature schemes in the GPV framework. Our main contributions are:

EXTENDING THE GPV FRAMEWORK TO RÉNYI DIVERGENCE. We extend the GPV framework to incorporate the Rényi divergence, adapting key lemmata to support the Rényi divergence and NTRU rings. These results are broadly applicable to other constructions including [EFG⁺22, ENS⁺23, GJK24, YJW23]. We also develop tools for optimally selecting parameters for Rényi divergence. While these contributions are not fundamentally new [SS11, LPR13, BLL⁺15, TT15], we present them here in full due to their practical significance. For instance, while FALCON recommends using a Rényi divergence of order $a = 2\lambda$, this results in a 60-bit security loss for the FALCON-1024 parameter set. Our tools reduce this loss to just 8 bits.

FALCON⁺: MODIFICATIONS TO FALCON FOR PROVABLE SECURITY. While our extensions to the GPV framework and parameter optimisation tools improve the security analysis, we were not able to prove the security of FALCON without modifications. To this end, we introduce FALCON⁺, a minor modification of FALCON, that can easily be justified at this late stage of the standardisation process. The differences to FALCON are sketched in Figure 1. Besides hashing the public key (which is standard cryptographic practice), FALCON⁺ crucially samples a random salt and samples a preimage of the hash of the message/salt pair *within* the repeat loop of signing⁵, i.e., until a sufficiently short preimage is found. In contrast, FALCON picks a fixed random salt *outside* of the repeat loop and then samples the preimage.⁶ This modification incurs minimal additional cost since the loop is executed only once or twice in expectation. Furthermore, the costs associated with Gaussian sampling within the loop far outweigh the hashing and FFT costs, even for large messages. Our proposed changes have already been integrated into the latest implementation of FALCON [Por25a, Por25b] and are due to be integrated into the forthcoming FIPS standard.

<u>Sgn(sk, m)</u>	<u>Sgn⁺(sk, m)</u>
01 Sample salt r	06 repeat
02 repeat	07 Sample salt r
03 $s \xleftarrow{\$} f^{-1}(H(r, m))$	08 $s \xleftarrow{\$} f^{-1}(H(pk, r, m))$
04 until $\ s\ _2 \leq \beta$	09 until $\ s\ _2 \leq \beta$
05 $\sigma := (r, s)$	10 $\sigma := (r, s)$

Figure 1. Signing (simplified) of original FALCON (left) and our **modification** FALCON⁺ (right). Sampling from $f^{-1}(\cdot)$ uses sk .

SECURITY ANALYSIS. We provide a thorough security analysis of FALCON⁺ in the random oracle model. Using our tools, we derive concrete security bounds from our theorems, which focus on minimising bit security loss due to Rényi divergence arguments. We formalise the hardness assumptions that are not only sufficient to prove the security of FALCON⁺ but also necessary, thereby providing clear targets for cryptanalysts. Specifically, the *plain unforgeability* of FALCON⁺ is equivalent to a multi-target inhomogeneous SIS problem, while *strong unforgeability* corresponds to a second-preimage version of the same problem, which we define here. Assuming that both problems are as hard as standard SIS, we show that FALCON⁺-512 (NIST Level I) achieves 113 bits of provable security for both *plain unforgeability* and *strong unforgeability*. Furthermore, by reducing the number of allowed signing queries from 2^{64} to 2^{58} , this increases to 119 bits, nearing the claimed security level. For FALCON⁺-1024 (NIST Level V), we prove that it meets 256 bits of security for both *plain unforgeability* and *strong unforgeability*. An overview of the provable bit security is shown in Table 1.

FFO SAMPLER. The Fast-Fourier Orthogonalization (FFO) process, introduced by Ducas and Prest [DP16], improves the running time of matrix orthogonalisation for matrices with circulant blocks. When the matrix dimensions are powers of 2, the time complexity improves by a factor of $\mathcal{O}((n/\log n))$, where n is the block size. This can be seen as a structured variant of the Gram–Schmidt algorithm. In FALCON, the FFO algorithm accelerates Gaussian sampling through the FFO SAMPLER. A proof of the FFO SAMPLER has not been published, neither in the FALCON specification nor in [Pre17], where the GPV sampler is analysed using the Rényi divergence rather than statistical distance. This use of the Rényi divergence, as opposed to statistical distance, enables the reduction of floating-point precision to 53 bits while maintaining a 256-bit security level. In Section F.3, we show that a similar result to the one from [Pre17] also holds for the FFO SAMPLER.

⁵ Actually, the message will be hashed once outside the repeat loop, and the resulting digest will then be hashed together with a fresh salt inside the repeat loop.

⁶ Note that SQUIRRELS [ENST23], a scheme submitted to the first round of the *NIST Call for Additional Post-Quantum Signature Schemes*, suffers from the same shortcoming.

Scheme	Notion	Multiplicative Loss and Assumption	Bit Security
FALCON ⁺ -512 ($Q_s = 2^{64}$)	UF-CMA (Th. 1) SUF-CMA (Th. 2)	$r_u^{Q_s} \cdot r_p^{Q_s} \cdot Q_H \cdot \mathcal{R}\text{-ISIS}_\beta$ UF-CMA + $r_p^{Q_s} \cdot Q_s \cdot \mathcal{R}\text{-SPISIS}_\beta$	113
FALCON ⁺ -512 ($Q_s = 2^{58}$)			119
FALCON ⁺ -1024 ($Q_s = 2^{64}$)			256

Table 1. Provable bit security levels of FALCON⁺-512 and FALCON⁺-1024, along with the simplified concrete security loss for FALCON⁺ in the random oracle model. Constants $r_u = 1 + \delta_u$ and $r_p = 1 + \delta_p$ represent Rényi divergences related to the uniformity of an NTRU evaluation on Gaussian inputs (r_u) and the preimage sampler (r_p). Q_s and Q_H denote the number of signing and random oracle queries, respectively.

This strengthens the theoretical foundation of the FALCON signature scheme and provides formal backing for its security claims.

1.2 Technical Overview

THE GENTRY-PEIKERT-VAIKUNTANATHAN FRAMEWORK. The GPV framework [GPV08] provides a method for constructing secure lattice-based signature schemes using the full-domain-hash (FDH) paradigm [BR96], commonly referred to as “*hash-and-sign*”. Central to this framework is a “*preimage sampleable trapdoor function*” a primitive, instantiated (in part) by a function $f_{\mathbf{A}}(s) := \mathbf{A}s \bmod q$ where $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$. Here, each signature essentially corresponds to a short preimage of the hash of a message. More specifically, the public key pk is a statistically close to a uniform matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ (with $n \leq m$) of the orthogonal q -ary “SIS” lattice $\Lambda = \{z \in \mathbb{Z}^m \mid \mathbf{A}z = 0 \bmod q\}$. The secret key (or trapdoor) sk is a matrix $\mathbf{B} \in \mathbb{Z}^{m \times m}$ orthogonal to $\mathbf{A}$, i.e., $\mathbf{A} \cdot \mathbf{B} = \mathbf{0}$. Provided the Gram–Schmidt norm of $\mathbf{B}$ is small, a short preimage under $f_{\mathbf{A}}$ can be found efficiently using $\mathbf{B}$. A signature on a message m is a short vector $s \in \mathbb{Z}^m$ such that $\mathbf{H}(m) = \mathbf{A}s \bmod q$, where $\mathbf{H} : \{0, 1\}^* \rightarrow \mathbb{Z}_q^n$ is a hash function. Verification involves checking both the shortness of s and that $f_{\mathbf{A}}(s) = \mathbf{H}(m)$. We consider the probabilistic (or salted) variant of the scheme, where a short preimage of $\mathbf{H}(m, r)$, for a random salt r , forms part of the signature.

THE GPV PROOF TEMPLATE. The GPV framework was proven secure in both the random oracle model [BR93, GPV08] and the quantum random oracle model [BDF⁺11] under the plain (unstructured) SIS assumption [Ajt96]. Security can be established in two ways: (1) via *collision resistance* of $f_{\mathbf{A}}$, reducing to SIS, or (2) via *one-wayness* of $f_{\mathbf{A}}$, reducing to ISIS. The original work [GPV08] provided a tight proof of *strong unforgeability* for FDH, leveraging collision resistance. In this overview, we focus on the one-wayness proof.

Suppose, for the sake of contradiction, that an adversary $\mathbf{A}$ breaks the *plain unforgeability* of the signature scheme, producing a forgery s^* for a message m^* and salt r^* , where s^* is short and $\mathbf{H}(m^*, r^*) = \mathbf{A}s^* \bmod q$. We construct a reduction $\mathbf{B}$ that solves the one-wayness of $f_{\mathbf{A}}$ on image u by using $\mathbf{A}$ as a subroutine. The reduction proceeds as follows:

- Set the public key pk of the signature scheme to be the matrix $\mathbf{A}$ from the one wayness game.
- Whenever $\mathbf{A}$ makes a signing query on message m , the random oracle is programmed for each fresh query to $\mathbf{H}(m, r)$. The reduction samples a Gaussian vector s_m , programs $\mathbf{H}(m, r) := \mathbf{A}s_m \bmod q$, and returns the signature (s_m, r) to $\mathbf{A}$. Crucially, by the “*leftover hash lemma*” [HILL99], the simulated random oracle output is statistically close to uniform.
- Program the hash of the target message m^* to be the one-wayness target vector, $\mathbf{H}(m^*, r^*) := u$.⁷
- When $\mathbf{A}$ outputs a forgery (s^*, r^*) for m^* , the reduction outputs s^* as a solution to the one-wayness challenge. By construction, it holds $f_{\mathbf{A}}(s^*) = \mathbf{H}(m^*, r^*) = u$, and s^* is short, so $\mathbf{B}$ succeeds.

Clearly, the one-wayness of $f_{\mathbf{A}}$ with target u can be directly reduced to an ISIS instance on input $(\mathbf{A}, u)$.

⁷ For simplicity in this overview, we ignore losses due to guessing and do not address the multi-target assumption.

FALCON INSTANTIATION OF THE GPV FRAMEWORK. The design of FALCON prioritises compactness, minimising the combined size of $|pk| + |\sigma|$. To achieve this, FALCON relies on the class of NTRU lattices introduced by Hoffstein, Pipher, and Silverman [HPS98, HHP⁺03], which come with an additional ring structure that reduces the public key size by a factor of $\mathcal{O}(n)$ and accelerates many computations by a factor of at least $\mathcal{O}(n/\log n)$. Among structured lattices, NTRU lattices are particularly efficient, with public keys represented as a single polynomial $\mathbf{h} \in \mathcal{R}_q = \mathbb{Z}_q[X]/(X^n + 1)$. FALCON instantiates the randomised version of the GPV framework, where the matrix $\mathbf{A}$ has the form $\mathbf{A} = [\mathbf{1} \ \mathbf{h}]$. Thus, the NTRU-based preimage sampleable trapdoor function $f_{\mathbf{h}}$ is simply the GPV function instantiated with this structured choice of $\mathbf{A}$ [HPS98, DLP14, PFH⁺22]. Specifically, $f_{\mathbf{h}}$ maps two ring elements $(\mathbf{s}_1, \mathbf{s}_2)$ to $\mathbf{s}_1 + \mathbf{h} \cdot \mathbf{s}_2 \bmod q$. Observe that $f_{\mathbf{h}}$ is a special case of the GPV trapdoor function $f_{\mathbf{A}}(s) = \mathbf{A}s \bmod q$. A valid signature on message m consists of a tuple $(\mathbf{s}_1, \mathbf{s}_2) \in \mathcal{R}^2$ and a random salt $r \in \{0, 1\}^k$ satisfying

$$H(m, r) = \mathbf{s}_1 + \mathbf{h} \cdot \mathbf{s}_2 \bmod q \quad \wedge \quad \|(\mathbf{s}_1, \mathbf{s}_2)\|_2 \leq \beta.$$

This adaptation requires the standard “*randomised GPV*” proof to be based on an “*NTRU-SIS*” assumption, a process described as “*straightforward*” in the FALCON specification [PFH⁺22].

REPEATED SAMPLING AND SALTING. One key difference between FALCON and the GPV framework is how large-norm signatures are handled. In GPV, the sampler may output signatures whose norm exceeds the required bound with some probability $\exp(-\Omega(n))$. These events are negligible but not explicitly filtered. In FALCON, signatures are not output directly from the sampler. Any signature whose ℓ_2 -norm exceeds the verification bound is rejected, and sampling is repeated. The rejection probability is at most 2^{-14} . Experiments indicate that it is even smaller in practice [Por25c]: about $2^{-17.8}$ for FALCON-512 and 2^{-30} for FALCON-1024.⁸ This rejection step introduces complications for simulating signing queries, as the process involves conditional distributions. The signing oracle outputs preimages conditioned on having a sufficiently small norm, whereas programming the random oracle with this constraint and analysing the uniformity of outputs appears to be challenging.

In the current FALCON specification, the random salt r is chosen before the preimage sampling loop and therefore does not help mitigate the issue of conditional distributions. In our modified scheme, FALCON⁺, we propose drawing a new salt each time the preimage sampling process results in too large signatures. This modification allows the reduction to continue programming the random oracle with large preimages, while still being able to produce valid signatures. If a sampled preimage is too large, the reduction can simply choose a new salt, yielding a new random oracle output and a new preimage. This change incurs only a minor constant overhead in the security bound, corresponding to the maximum number of repetitions. In practice, the efficiency impact is minimal, as preimage sampling remains the dominant computational cost in both the original and modified schemes. The latest FALCON implementation incorporates these changes [Por25a, Por25b], and the forthcoming FIPS standard will include them as well.

RÉNYI DIVERGENCE IN FALCON. Another issue is that FALCON relies on the Rényi divergence, whereas the GPV framework uses the statistical (or total variation) distance to prove the closeness of the sampler and a Gaussian. Citing [Pre17, Lem. 6] as the analysis of the KLEIN SAMPLER [Kle00], FALCON claims that for suitable parameters, the Rényi divergence between the FFO SAMPLER’s output and an ideal Gaussian is bounded by $1 + \mathcal{O}(1)/Q_s$, incurring a loss of at most $\mathcal{O}(1)$ bits of security. However, we are interested in the concrete bounds. To address this, we modify the GPV framework to handle Rényi divergence, enabling the simulation of signing queries.

Furthermore, the GPV framework uses a second statistical argument, the “*leftover hash lemma*” [HILL99, Lem. 4.8], to show that the programmed output of the random oracle is close to uniform. However, two challenges arise. First, the argument, originally stated for unstructured lattices, must be adapted to the ring setting, which can be done using a regularity lemma from [SS11, Sec. 3.3] or [LPR13, Sec. 4]. More critically, applying such a statistical argument to the FALCON parameters yields statistical distances as large

⁸ It may also happen that the norm check passes, but compression fails. This is about 25 times less likely for the small parameter set, and about 10^{-3} for the large one. If NIST were to enforce an ℓ_∞ -norm check instead, the restart probability would increase to about 2^{-10} , making salt resampling even more important.

as 2^{-34} , for each simulated random oracle output. As a result, further modifications to the GPV framework are necessary to argue that the random oracle’s output is Rényi-close to uniform. That is, we require a lemma showing that $H(m, r) := \mathbf{s}_1 + \mathbf{h} \cdot \mathbf{s}_2 \bmod q$ is Rényi close to uniform for Gaussian $\mathbf{s}_1, \mathbf{s}_2$. However, the Rényi divergence arguments are highly sensitive to the number of queries, and the FALCON parameters are specifically tuned to accommodate the number of signing queries, $Q_s = 2^{64}$, rather than the random oracle queries, $Q_H = 2^{96} \gg Q_s$. Thus, these tools cannot be applied directly in the random oracle model, requiring us to carefully program only those random oracle queries originating from signing queries.

NORM BOUND. GPV [GPV08] showed that *strong unforgeability* follows from the collision resistance of f_A , which reduces to SIS with norm bound 2β . Similarly, *plain unforgeability* follows from the one-wayness of f_A , reducing to ISIS with norm bound only β . While the plain unforgeability proof extends to FALCON, the strong unforgeability proof does not: under FALCON-1024 parameters, SIS with norm 2β is trivial, and for FALCON-512, it falls short of the desired security margin. This motivates a proof of *strong unforgeability* where the resulting (I)SIS instance has norm bound β . We show that such a proof exists under the second preimage resistance of f_A , which reduces to a corresponding second preimage variant of the ISIS problem with norm bound β . Refer to Figure 2 for an overview.




Figure 2. Relationships between computational hardness assumptions, properties of the preimage-sampleable trapdoor function f_A , and security notions of the resulting signature scheme. Our results are highlighted by the dotted box.

FFO SAMPLER. The FALCON specification does not explicitly analyse the FFO SAMPLER; instead, it bounds the relative error by relying on an analysis of the KLEIN SAMPLER presented in [Pre17, Sec. 4.5]. The technique used in this analysis consists of proving the probability that the KLEIN SAMPLER will output the vector $\mathbf{z}$. This distribution is then compared to the discrete Gaussian distribution using Rényi divergence tools. Our analysis of the FFO SAMPLER follows the same approach, with the key step being the proof of its distribution. The FFO SAMPLER works with a matrix of polynomials, whereas the KLEIN SAMPLER operates on a matrix in $\mathbb{R}^{2n \times 2n}$. The KLEIN SAMPLER relies on two fundamental results. The first states that in a lattice generated by an orthogonal basis, one can sample vectors from a Gaussian distribution by independently sampling each coordinate from a discrete Gaussian over the integers. The second result provides a transformation that reduces the problem of sampling from a lattice generated by an arbitrary basis $\mathbf{B}$ to the problem of sampling from the lattice defined by its Gram–Schmidt orthogonalisation.

The FFO SAMPLER builds on the same principles as the KLEIN SAMPLER, but also leverages the tower-of-fields $\mathbb{K}_n/\mathbb{K}_{n/2}/\dots/\mathbb{K}_2/\mathbb{Q}$, where $\mathbb{K}_i = \mathbb{Q}[x]/(x^i + 1)$, to reduce complexity via a recursive algorithm. The tower-of-fields is navigated using linear mappings V and M . In particular, the mapping V is an isometry, which ensures that the Gaussian distribution of the output vector is preserved as an invariant throughout the recursion. To sample from the lattice generated by the basis $\mathbf{B} = (\mathbf{b}_1, \mathbf{b}_2)$ around a target vector $(\mathbf{t}_1, \mathbf{t}_2)$, the FFO SAMPLER performs $2 \log n$ recursive calls. At the top of the recursion tree, a transformation (Lemma 18) reduces the problem to sampling in the lattice generated by the Gram–Schmidt orthogonalized basis $\tilde{\mathbf{B}} = (\tilde{\mathbf{b}}_1, \tilde{\mathbf{b}}_2)$ around the vector $(\mathbf{t}'_1, \mathbf{t}'_2)$. Since $\tilde{\mathbf{B}}$ is orthogonal, the two coordinates can be sampled independently (Lemma 17). The vector $\tilde{\mathbf{b}}_1$ is transformed into a 2×2 matrix $(\mathbf{b}_{11}, \mathbf{b}_{12})$ whose entries are polynomials of

degree $n/2$ using a linear mapping M . Similarly, $\mathbf{t}'_1$ is mapped to a pair $(\mathbf{t}_{11}, \mathbf{t}_{12})$, consisting of polynomials of degree $n/2$, using a mapping V . These inputs are then recursively sampled to yield a Gaussian vector $(\mathbf{z}_{11}, \mathbf{z}_{12})$. This allows us to recover the coordinates $\mathbf{z}_1 = V^{-1}(\mathbf{z}_{11}, \mathbf{z}_{12})$, which itself follows a Gaussian distribution. The same procedure is applied to obtain $\mathbf{z}_2$. At each step of the recursion, the number of columns of the input matrix remains two, while the number of rows doubles. At the leaves of the recursive tree, the entries of the matrix are reals (polynomials of degree 0), and since the number of columns is 2, the KLEIN SAMPLER can be applied directly in dimension 2. The analysis of the FFO SAMPLER builds upon that of the KLEIN SAMPLER, leveraging the isometry property of the mapping V .

CONCRETE SECURITY OF FALCON⁺. Table 1 summarises the concrete security bounds of FALCON⁺, our modified version of FALCON. Here, t - $\mathcal{R}$ -ISIS denotes the multi-target variant of the ring inhomogeneous SIS problem with t targets, while t - $\mathcal{R}$ -SPISIS denotes a second-preimage variant, where an adversary is given t targets with valid short preimages and must output a different short preimage for one of them. The FALCON parameters have been carefully chosen so that the Rényi divergence bound $r_u^{Q_s} = (1 + \delta_u)^{Q_s}$ remains a small constant for $Q_s = 2^{64}$ signing queries. This allows signing queries to be simulated by programming the random oracle. The forgery in the (plain) unforgeability game is related to one of Q_H direct random oracle queries. To use the forgery, the proof of **UF-CMA** in Theorem 1 relies on the t - $\mathcal{R}$ -ISIS assumption with $t \approx Q_H$ targets. The challenge targets are embedded into the random oracle, allowing one to be solved upon receiving a forgery. The Rényi loss introduced is $r_u^{Q_s}$ due to the changes in the signing distribution. For FALCON⁺-512, parameters were selected with a narrow margin, resulting in a Rényi loss of 7 bits when $Q_s = 2^{64}$. Reducing the allowed number of signing queries to $Q_s = 2^{58}$ lowers the loss to just 1 bit. For FALCON⁺-1024, the Rényi loss is 8 bits when $Q_s = 2^{64}$ signing queries, but due to the substantial security margin for larger SIS instances, this does not compromise the target 256-bit security.

Theorem 2 for *strong unforgeability* reduces to both **UF-CMA** and the second-preimage version t - $\mathcal{R}$ -SPISIS. Here, the short preimages are used to simulate signing queries, and the *strong* forgery is used to obtain a second preimage. Using techniques from [BBD⁺23, FFH25], we can set $t = Q_s$. Assuming t - $\mathcal{R}$ -SPISIS is as hard as standard SIS, and factoring in Rényi losses, the dominant term in Theorem 2 is the **UF-CMA** term. Thus, the concrete security of **SUF-CMA** and **UF-CMA** is essentially the same for both parameter sets.

The resulting bit security levels for FALCON⁺-512 (NIST Level I) and FALCON⁺-1024 (NIST Level V) are shown in Table 1. These values are derived from Theorem 1 and Theorem 2, taking into account the Rényi loss for the FALCON parameter sets, and using the “lattice-estimator” [APS15a, APS15b] to estimate the hardness of SIS.

Open Problems. Finally, we leave as an open problem a proof in the quantum random oracle model (QROM), which could likely be achieved using the techniques from [BBD⁺23, FFH25], provided that the Rényi arguments can be handled correctly.

2 Preliminaries

We introduce some relevant notation and definitions used throughout the paper.

2.1 Notation

SETS AND ALGORITHMS. We write $s \xleftarrow{\$} \mathcal{S}$ to denote the uniform sampling of s from the finite set $\mathcal{S}$ and by $\mathcal{U}(\mathcal{S})$ the uniform distribution over $\mathcal{S}$. For an integer n , we define $[n] := \{1, \dots, n\}$. The notation $\llbracket b \rrbracket$, where b is a boolean statement, evaluates to 1 if the statement is true and 0 otherwise. We use uppercase letters A, B, C, D to denote algorithms. Unless otherwise stated, algorithms are probabilistic, and we write $(y_1, \dots) \xleftarrow{\$} A(x_1, \dots)$ to denote that A returns $(y_1, \dots)$ when run on input $(x_1, \dots)$. We write A^B to denote that A has oracle access to B during its execution. The support of a discrete random variable X is defined as $\text{supp}(X) := \{x \in \mathbb{R} \mid \Pr[X = x] > 0\}$. For two polynomials $\mathbf{f}, \mathbf{g} \in \mathcal{R}_q = \mathbb{Z}_q[X]/(X^n + 1)$, we denote the

polynomial multiplication of $\mathbf{f}$ and $\mathbf{g}$ by $\mathbf{f} \cdot \mathbf{g}$. When the rank needs to be made explicit, we write $\mathcal{R}_q(n)$. We use $\bar{\mathcal{R}}_q(n)$ to denote polynomials in $(\mathbb{R}/q\mathbb{Z})[X]/(X^n + 1)$. By “log” we denote the logarithm of base 2, by “ln” of base e . We use $\lesssim$ to denote an approximate inequality.

SECURITY GAMES. We use standard code-based security games [BR06]. A *game* G is a probability experiment in which an adversary A interacts with an implicit challenger that answers oracle queries issued by A . The game G has one *main procedure* and an arbitrary amount of additional *oracle procedures* which describe how these oracle queries are answered. We denote the (binary) output b of game G between a challenger and an adversary A as $G(A) \Rightarrow b$. A is said to *win* G if $G(A) \Rightarrow 1$, or shortly $G \Rightarrow 1$. Unless otherwise stated, the randomness in the probability term $\Pr[G(A) \Rightarrow 1]$ is over all the random coins in game G and adversary A . To provide a cleaner description and avoid repetitions, we sometimes refer to procedures of different games. To call the oracle procedure Oracle of game G on input x , we shortly write $G.\text{Oracle}(x)$. If a game is aborted the output is 0. For our analysis we rely on the commonly used main difference lemma or the multiplicative difference lemma for independent events. Security notions are considered in the random oracle model [BR93].

2.2 Signatures

We recall the syntax and standard security notions of signatures.

Definition 1 (Signature Scheme). A *signature scheme* Sig is defined as a tuple $(\text{Gen}, \text{Sgn}, \text{Ver})$ of the following three algorithms.

- $(sk, pk) \xleftarrow{\$} \text{Gen}$: The probabilistic key generation algorithm returns a secret key sk and a corresponding public key pk , where pk defines a message space $\mathcal{M}$.
- $\sigma \xleftarrow{\$} \text{Sgn}(sk, m)$: Given a secret key sk and a message $m \in \mathcal{M}$, the probabilistic signing algorithm Sgn returns a signature σ .
- $b \leftarrow \text{Ver}(pk, m, \sigma)$: Given a public key pk , a message m , and a signature σ , the deterministic verification algorithm Ver returns a bit b , such that $b = 1$ if and only if σ is a valid signature on m and $b = 0$ otherwise.

Sig has ε -*correctness error* if for all $(sk, pk) \in \text{sup}(\text{Gen})$ and any $m \in \mathcal{M}$ $\Pr[\text{Ver}(pk, m, \text{Sgn}(sk, m)) \neq 1] \leq \varepsilon$, where the probability is taken over the random choices of Sgn .

Definition 2 ((Strong) Unforgeability). The notions of *(strong) existential unforgeability under chosen message attacks* are formalised via the games $Q_s\text{-UF-CMA}_{\text{Sig}}(A)$ and $Q_s\text{-SUF-CMA}_{\text{Sig}}(A)$. Both are depicted in Figure 3, where Q_s is an upper bound on the adversary’s signing queries. We define the advantage functions of adversary A as

$$\begin{aligned} \text{Adv}_{\text{Sig}, A}^{Q_s\text{-UF-CMA}} &:= \Pr[Q_s\text{-UF-CMA}_{\text{Sig}}(A) \Rightarrow 1], \\ \text{Adv}_{\text{Sig}, A}^{Q_s\text{-SUF-CMA}} &:= \Pr[Q_s\text{-SUF-CMA}_{\text{Sig}}(A) \Rightarrow 1]. \end{aligned}$$

2.3 Lattices

RINGS AND NORMS. In this work, we work with polynomial rings of the form $\mathcal{R} = \mathbb{Z}[X]/(X^n + 1)$ and $\mathcal{R}_q = \mathbb{Z}_q[X]/(X^n + 1)$, for $n = 2^k$ and $k, q \in \mathbb{N}$. For a polynomial $\mathbf{f} \in \mathcal{R}_q$, let $f \in \mathbb{Z}_q^n$ denote the coefficient embedding of $\mathbf{f}$, and $f_i \in \mathbb{Z}_q$ the i^{th} coefficient.

Definition 3 (Anticirculant Matrix). For a polynomial $\mathbf{f} \in \mathcal{R} = \mathbb{Z}[X]/(X^n + 1)$, the anticirculant matrix of $\mathbf{f}$ is defined as

$$\mathcal{A}(\mathbf{f}) = \begin{bmatrix} f_0 & -f_{n-1} & \cdots & -f_1 \\ f_1 & f_0 & \cdots & -f_2 \\ \vdots & \vdots & \ddots & \vdots \\ f_{n-1} & f_{n-2} & \cdots & f_0 \end{bmatrix} \in \mathbb{Z}^{n \times n}.$$

Games $Q_s\text{-UF-CMA}_{\text{Sig}}(\mathbf{A})/Q_s\text{-SUF-CMA}_{\text{Sig}}(\mathbf{A})$	
01	$\mathcal{Q} \leftarrow \emptyset$
02	$(sk, pk) \xleftarrow{\$} \text{Gen}$
03	$(m^*, \sigma^*) \xleftarrow{\$} \mathbf{A}^{\text{Sgn}(\cdot)}(pk)$
04	return $\llbracket \text{Ver}(pk, m^*, \sigma^*) = 1 \wedge (m^*, \cdot) \notin \mathcal{Q} \rrbracket$ // UF-CMA
05	return $\llbracket \text{Ver}(pk, m^*, \sigma^*) = 1 \wedge (m^*, \sigma^*) \notin \mathcal{Q} \rrbracket$ // SUF-CMA
Oracle $\text{Sgn}(m)$	
06	$\sigma \xleftarrow{\$} \text{Sgn}(sk, m)$
07	$\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(m, \sigma)\}$
08	return σ

Figure 3. Games defining **UF-CMA** and **SUF-CMA** for a signature scheme $\text{Sig} = (\text{Gen}, \text{Sgn}, \text{Ver})$ and adversary $\mathbf{A}$ making at most Q_s queries to Sgn .

Anticirculant matrices satisfy the following useful properties.

Lemma 1. *Let $\mathbf{f}, \mathbf{g} \in \mathcal{R}$. Then $\mathcal{A}(\mathbf{f}) + \mathcal{A}(\mathbf{g}) = \mathcal{A}(\mathbf{f} + \mathbf{g})$ and $\mathcal{A}(\mathbf{f}) \cdot \mathcal{A}(\mathbf{g}) = \mathcal{A}(\mathbf{f} \cdot \mathbf{g})$.*

This implies an isomorphism between $\mathcal{R}$ and the anticirculant matrices over $\mathbb{Z}^{n \times n}$, $\mathcal{R}_q$ and $\mathbb{Z}_q^{n \times n}$ respectively. Sometimes we overload the notation and write $\mathcal{A}(f)$ for the coefficient embedding $f \in \mathbb{Z}^n$ of $\mathbf{f}$ instead of $\mathcal{A}(\mathbf{f})$.

Let the ℓ_2 -norm for $\mathbf{f} = f_0 + f_1X + \dots + f_{n-1}X^{n-1} \in \mathcal{R}$ be defined as $\|\mathbf{f}\|_2 := \sqrt{\sum_{i=0}^{n-1} |f_i|^2}$. For two polynomials $\mathbf{f}, \mathbf{g} \in \mathcal{R}$ we use the notation

$$\|(\mathbf{f}, \mathbf{g})\|_2 := \sqrt{\sum_{i=0}^{n-1} (|f_i|^2 + |g_i|^2)}.$$

LATTICES. A lattice $\mathbf{\Lambda} \subseteq \mathbb{R}^n$ is a discrete additive subgroup of $\mathbb{R}^n$.

Definition 4 (Lattice). A rank m lattice in $\mathbb{R}^n$ is defined via the set $b_1, \dots, b_m \in \mathbb{R}^n$ of *linearly independent* vectors that form a basis $\mathbf{B} = \{b_1, \dots, b_m\}$ for the lattice

$$\mathbf{\Lambda} := \mathbf{\Lambda}(\mathbf{B}) = \mathbf{\Lambda}(b_1, \dots, b_m) = \left\{ \sum_{i=1}^m c_i b_i \mid c_1, \dots, c_m \in \mathbb{Z} \right\}.$$

If $m = n$, then $\mathbf{\Lambda}$ is a full-rank lattice.

The *determinant* of a lattice $\mathbf{\Lambda} = \mathbf{\Lambda}(\mathbf{B}) \subseteq \mathbb{R}^n$ for some basis $\mathbf{B} \in \mathbb{R}^{n \times m}$ is defined as $\det(\mathbf{\Lambda}) = \sqrt{\det(\mathbf{B}^\top \mathbf{B})}$. For an n -dimensional lattice $\mathbf{\Lambda}$, a lattice $\mathbf{\Lambda}' \subseteq \mathbf{\Lambda}$ is called a sublattice of $\mathbf{\Lambda}$. The *shifted lattice* by $t \in \mathbb{R}^n$ is denoted by $\mathbf{\Lambda} + t = \{x + t \mid x \in \mathbf{\Lambda}\}$. One can define the following quotient group $\mathbf{\Lambda}/\mathbf{\Lambda}' := \{t + \mathbf{\Lambda}' \mid t \in \mathbf{\Lambda}\}$, which forms a group under the addition of cosets $t + \mathbf{\Lambda}'$. The *orthogonal* lattice for $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ is defined as $\mathbf{\Lambda}^\perp(\mathbf{A}) := \{e \in \mathbb{Z}^m \mid \mathbf{A}e = \mathbf{0} \pmod{q}\}$ and its *shifted lattice*, for a shift $t \in \mathbb{Z}^n$, is defined as $\mathbf{\Lambda}_t^\perp(\mathbf{A}) := \{e \in \mathbb{Z}^m \mid \mathbf{A}e = t \pmod{q}\}$. If $\mathbf{\Lambda}$ is an *orthogonal* lattice, then $\mathbf{\Lambda}_t$ denotes its shift by t .

Definition 5 (NTRU Lattice). Let $n = 2^k$ for $k \in \mathbb{Z}$, q prime, $\mathbf{f}, \mathbf{g} \in \mathcal{R} = \mathbb{Z}[X]/(X^n + 1)$, and $\mathbf{h} = \mathbf{g} \cdot \mathbf{f}^{-1} \pmod{q}$. The NTRU lattice parameterised by $\mathbf{h}$ and q is a lattice of volume q^n in $\mathbb{R}^{2n}$ in the coefficient embedding of the following module

$$\{(\mathbf{u}, \mathbf{v}) \in \mathcal{R}^2 \mid \mathbf{u} + \mathbf{v} \cdot \mathbf{h} = \mathbf{0} \pmod{q}\}.$$

Sometimes, we use the notation $\mathbf{\Lambda}_{\mathbf{h}, q} = \mathbf{\Lambda}^\perp(\mathbf{A})$ for $\mathbf{A} = [I_N \mid \mathcal{A}(\mathbf{h})] \in \mathbb{Z}^{n \times 2n}$.

Equivalently, for $\mathcal{R} = \mathbb{Z}[X]/(X^n + 1)$, an NTRU lattice is a full-rank submodule lattice of $\mathcal{R}^2$ generated by the columns of a matrix of the form

$$B_h = \begin{bmatrix} -h & q \\ 1 & 0 \end{bmatrix}$$

for prime q , $q = q \cdot 1$, and some $h \in \mathcal{R}_q$. A trapdoor for this lattice is a relatively short basis

$$B_{f,g} = \begin{bmatrix} g & G \\ -f & -F \end{bmatrix}$$

where the basis vectors $(f, g) \in \mathcal{R}^2$ and $(F, G) \in \mathcal{R}^2$ are not much larger than $\sqrt{\det B_h} = \sqrt{q}$ and $f \cdot G - g \cdot F = q \pmod{X^n + 1}$.

GAUSSIANS AND PREIMAGE SAMPLING. We define discrete Gaussians and state some of their useful properties.

Definition 6 (Discrete Gaussian Distribution over $\Lambda + t$). The n -dimensional *Gaussian function* $\rho_{s,c}: \mathbb{R}^n \rightarrow (0, 1]$ on $\mathbb{R}^n$ centered at $c \in \mathbb{R}^n$ with width $s > 0$ is defined by

$$\rho_{s,c}(x) := \exp\left(-\frac{\|x - c\|_2^2}{2s^2}\right).$$

For any $c \in \mathbb{R}^n$, $t \in \mathbb{R}^n$, $s \in \mathbb{R}^+$, and lattice Λ , the *discrete Gaussian distribution over $\Lambda + t$* is defined as

$$\forall x \in \Lambda + t, \quad \mathcal{D}_{\Lambda+t,s,c} := \frac{\rho_{s,c}(x)}{\sum_{z \in \Lambda+t} \rho_{s,c}(z)}.$$

We sometimes use the following notation $\rho_{s,c}(\Lambda + t) = \sum_{x \in \Lambda+t} \rho_{s,c}(x)$. We omit the subscript c when the Gaussian is centered at 0 and subscript Λ when the Gaussian is over $\mathbb{Z}^n$. We use $f \sim \mathcal{D}_{\mathcal{R}}$ to denote the polynomial $f := \sum_{i=0}^{n-1} f_i X^i \pmod{X^n + 1}$ for $f \sim \mathcal{D}_{\mathbb{Z}^n}$.

For bounding the probability that a random variable deviates a long way from the mean, we will use the following tail bound from [Ban93, Lyu12, DRSD14, ADRS15].

Lemma 2 (Gaussian Tail Bound (unnormalised version of [ADRS15, Lem. 2]). For any lattice $\Lambda \subseteq \mathbb{R}^n$, width $s > 0$, shift $t \in \mathbb{R}^n$, and tailcut rate $\tau > 1$,

$$\Pr_{z \leftarrow \mathcal{D}_{\Lambda+t,s}} [\|z\|_2 > \tau s \sqrt{n}] \leq \frac{\rho_s(\Lambda)}{\rho_s(\Lambda + t)} \left(\sqrt{e^{1-\tau^2} \tau^2} \right)^n.$$

Definition 7 (Gram–Schmidt Norm [GPV08, DLP14]). For a finite basis $B = (b_i)_{i \in I}$, let $\tilde{B} = (\tilde{b}_i)_{i \in I}$ be its Gram–Schmidt orthogonalization. Then the Gram–Schmidt norm of B is the value $\|B\|_{GS} := \max_{i \in I} \|\tilde{b}_i\|$.

Lemma 3 (NTRU Trapdoor Generation [HPS98, Pre15]). For a ring $\mathcal{R}$, the NTRU trapdoor generation algorithm $\text{TpGen}(\alpha, q) \rightarrow (f, g, F, G, h)$ takes a target quality $\alpha \geq 1$ and a modulus q , and returns a public key $h \in \mathcal{R}_q \setminus \{0\}$ together with the corresponding trapdoor $(f, g, F, G) \in \mathcal{R}^4$, such that B_h and $B_{f,g}$ form a basis of the same lattice. Furthermore, $\|B_{f,g}\|_{GS} \leq \alpha \sqrt{q}$. When convenient, we write $(B, h) \in \text{TpGen}$ for short.

Let Λ be an n -dimensional lattice and $\epsilon > 0$, the (scaled) smoothing parameter $\eta_\epsilon(\Lambda)$ is the smallest $s > 0$ such that $\rho_{1/s}(\Lambda^* \setminus 0) \leq \epsilon$, where Λ^* denotes the dual lattice (the exact definition of the dual is not required for this work). We will use the following upper bound on the smoothing parameter.

Lemma 4 (Special Case of [MR07, Lem. 3.3]). For any $\epsilon \in (0, 1)$ it holds that

$$\eta_\epsilon(\mathbb{Z}^{2n}) \leq \frac{1}{\pi} \cdot \sqrt{\frac{\ln(4n(1+1/\epsilon))}{2}}.$$

The following lemma appears implicitly in [MR04, MR07].

Lemma 5 (Implicit in [MR07, Lem. 4.4]). *For any n -dimensional lattice $\mathbf{\Lambda}$, center $c \in \mathbb{R}^n$, and reals $0 < \epsilon < 1$, $s \geq \eta_\epsilon(\mathbf{\Lambda})$, we have*

$$\rho_{s,c}(\mathbf{\Lambda}) \in [1 - \epsilon, 1 + \epsilon] \cdot \frac{(\sqrt{2\pi} \cdot s)^n}{\det(\mathbf{\Lambda})}.$$

FFO SAMPLER. The mappings M and V are similar to the mapping $\mathcal{A}$ used to define anticirculant matrices, but they work step by step. While $\mathcal{A}$ expresses matrices of polynomials in $\mathcal{R}_q$ as block matrices with elements from $\mathbb{Z}_q$, V and M express matrices of polynomials in $\mathcal{R}_q = \mathbb{Z}_q[X]/(X^n + 1)$ as block matrices of anticirculant matrices. These matrices are polynomials in $\mathbb{Z}_q[X]/(X^{n/2} + 1)$. The idea behind the FFO SAMPLER is to recursively apply the operator $M_{k/(k/2)}$, working with polynomials of degree $k/2$ at each step. The operator $V_{k/(k/2)}$ splits degree- k polynomials into their even and odd components, each of degree $k/2$, similar to the decomposition used in the FFT.

Conversely, $V_{k/(k/2)}^{-1}$ recombines the two halves into a single degree- k polynomial. The sampling procedure begins with the matrix $\mathbf{B}$, and at each step, the factorisation LDL^* is computed, corresponding to the Gram–Schmidt orthogonalisation. This step may be skipped by using the FALCON tree, which essentially precomputes the Gram–Schmidt orthogonalisation. At the leaves of the recursion tree, the matrix $\mathbf{B}$ contains integer entries in two columns, and the KLEIN SAMPLER is called. At each other step, the FFO SAMPLER takes care of computing with orthogonal basis. The operator $M_{k/(k/2)}$ transforms the orthogonal vectors $\tilde{\mathbf{b}}_1$ and $\tilde{\mathbf{b}}_2$ into 4 mutually orthogonal vectors $\tilde{\mathbf{b}}_{11}, \tilde{\mathbf{b}}_{12}, \tilde{\mathbf{b}}_{21}, \tilde{\mathbf{b}}_{22}$. Thus, each recursive call to FFO SAMPLER receives a basis consisting of two orthogonal column vectors. At the leaves, the KLEIN SAMPLER performs integer sampling, with a standard deviation depending on the initial standard deviation s and the norms $\|\tilde{\mathbf{b}}_i\|$. The notation $\odot$ denotes the multiplication in the ring $\mathcal{R}_{q,k} = \mathbb{Z}[X]/(X^k + 1, q)$ with $k|n$ a power of two. Figure 4 describes the FFO SAMPLER.

FFOSampler $\left(\mathbf{B} = [\mathbf{b}_1, \mathbf{b}_2] \in \mathcal{R}_{q,k}^{2^{n-k+1} \times 2}, s \in \mathbb{R}, \mathbf{t} = (t_1, t_2) \in \mathcal{R}_{q,k}^{2^{n-k+1} \times 2}\right)$

```

01  $(\mathbf{L}, \tilde{\mathbf{B}}) \leftarrow \text{LDL}^*(\mathbf{B})$  so that  $\mathbf{B} = \mathbf{L}\tilde{\mathbf{B}}$  and  $\tilde{\mathbf{B}} = [\tilde{\mathbf{b}}_1, \tilde{\mathbf{b}}_2]$ 
02 if  $k = 0$ , KleinSampler $(\mathbf{B} = [\mathbf{b}_1, \mathbf{b}_2] \in \mathbb{Z}^{n \times 2}, s, \mathbf{t} = (t_1, t_2) \in (\mathbb{R}^n)^2)$ 
03 else // KleinSampler en dimension 2 with recursive calls
04    $z_2 \leftarrow^s V_{k/(k/2)}^{-1}(\text{FFOSampler}(M_{k/(k/2)}(\tilde{\mathbf{b}}_2), s, V_{k/(k/2)}(t_2)))$ 
05    $t'_1 = t_1 - (z_2 - t_2) \odot L_{1,2}$ 
06    $z_1 \leftarrow^s V_{k/(k/2)}^{-1}(\text{FFOSampler}(M_{k/(k/2)}(\tilde{\mathbf{b}}_1), s, V_{k/(k/2)}(t'_1)))$ 
07 return  $\mathbf{z} = (z_1, z_2)$ 

```

Figure 4. FFO SAMPLER.

2.4 Rényi Divergence

Definition 8 (Rényi Divergence [Rén61, BLL⁺15, Pre17]). Let $\mathcal{P}, \mathcal{Q}$ be two distributions such that $\text{sup}(\mathcal{P}) \subseteq \text{sup}(\mathcal{Q})$. For $a \in (1, \infty)$, we define the Rényi divergence of order a as

$$R_a(\mathcal{P} \parallel \mathcal{Q}) = \left(\sum_{x \in \text{sup}(\mathcal{P})} \frac{\mathcal{P}(x)^a}{\mathcal{Q}(x)^{a-1}} \right)^{\frac{1}{a-1}}.$$

In addition, we define the Rényi divergence of order $+\infty$ as

$$R_\infty(\mathcal{P}||\mathcal{Q}) = \max_{x \in \text{sup}(\mathcal{P})} \frac{\mathcal{P}(x)}{\mathcal{Q}(x)}.$$

Note that it is not symmetric and does not satisfy the triangle inequality. When the Rényi divergence is finite, which it will be for all our applications, we can think of it as a value $1 + \delta$ for $\delta \geq 0$. A smaller δ indicates that the distributions are closer.

The Rényi divergence satisfies several useful properties. A detailed overview can be found in Section A.

Definition 9 (Relative Error (implicit in [Pre17, Lem. 3])). Let $\mathcal{P}$ and $\mathcal{Q}$ be two distributions such that $\text{sup}(\mathcal{P}) = \text{sup}(\mathcal{Q})$. The relative error of $\mathcal{P}$ and $\mathcal{Q}$ is defined as

$$\delta_{RE}(\mathcal{P}, \mathcal{Q}) := \max_{x \in \text{sup}(\mathcal{P})} \frac{|\mathcal{P}(x) - \mathcal{Q}(x)|}{\mathcal{Q}(x)}.$$

The following lemma shows that the relative error can be used to bound the Rényi divergence. The original proof in [Pre17, Lem. 3] uses a Taylor expansion to approximate the function as $\delta \rightarrow 0$, denoted $\lesssim$. For simplicity, when applying the lemma we only use $\leq$.

Lemma 6 (Relative Error [Pre17, Lem. 3]). Let $\mathcal{P}, \mathcal{Q}$ be two distributions such that $\text{sup}(\mathcal{P}) = \text{sup}(\mathcal{Q})$ and $\delta_{RE} > 0$. Then for all $a \in (1, \infty)$

$$R_a(\mathcal{P} || \mathcal{Q}) \lesssim 1 + \frac{a\delta_{RE}^2}{2}.$$

The KLEIN SAMPLER [Kle00, GPV08] was analysed in [Pre17] with respect to its relative error and Rényi divergence. We analyse the FFO SAMPLER (Fast Fourier Orthogonalization) from [DP16] as used in FALCON in Section F and state the main results here.

Lemma 7 (Relative Error of FFO Sampler). Let n be a positive integer and $\epsilon \in (0, 1/4)$. Then the *relative error* of $\text{PreSmp} := \text{FFO SAMPLER}$ and the lattice $\mathbf{\Lambda} = \mathbf{\Lambda}(\mathbf{B})_{(\mathbf{c}, \mathbf{0})}$ for any basis $\mathbf{B} \in \mathbb{Z}^{2n \times 2n}$, standard deviation $s \geq \eta_\epsilon(\mathbb{Z}^{2n}) \cdot \|\mathbf{B}\|_{GS}$, and arbitrary syndrome $\mathbf{c} \in \mathcal{R}_q$ is bounded by

$$\delta_{RE}(\text{PreSmp}(\mathbf{B}, s, (\mathbf{c}, \mathbf{0})), \mathcal{D}_{\mathbf{\Lambda}, s}) \leq \left(\frac{1 + \epsilon/2n}{1 - \epsilon/2n} \right)^{2n} - 1 \approx 2\epsilon.$$

The proof can be found in Section F.3.

Corollary 1 (Rényi Divergence of FFO Sampler). Let n be a positive integer, $a \in (1, \infty)$, and $\epsilon \in (0, 1/4)$. Then for $\text{PreSmp} := \text{FFO SAMPLER}$ and the lattice $\mathbf{\Lambda} = \mathbf{\Lambda}(\mathbf{B})_{(\mathbf{c}, \mathbf{0})}$, for any basis $\mathbf{B} \in \mathbb{Z}^{2n \times 2n}$, $s \geq \eta_\epsilon(\mathbb{Z}^{2n}) \cdot \|\mathbf{B}\|_{GS}$, and arbitrary syndrome $\mathbf{c} \in \mathcal{R}_q$, the *Rényi divergence* is bounded by

$$R_a(\text{PreSmp}(\mathbf{B}, s, (\mathbf{c}, \mathbf{0})) || \mathcal{D}_{\mathbf{\Lambda}, s}) \lesssim 1 + 2a\epsilon^2.$$

2.5 Hardness Assumptions

We define two inhomogeneous variants of the Short Integer Solution problem over NTRU lattices. The first is a *multi-target* version, where the adversary is given t challenges and may solve one of them. The second is a *second-preimage* version, where the adversary is given t targets together with valid short preimages⁹ and must output a distinct short preimage for one of them.

⁹ Note that the game itself does not need to run in polynomial time; only the adversary is required to be efficient.

Definition 10 (t - $\mathcal{R}$ -ISIS, t - $\mathcal{R}$ -SPISIS). Let $t \geq 1$ and $\mathcal{R}$ be a ring. The *Ring Inhomogeneous Short Integer Solution* problem and the *Ring Second-preimage Inhomogeneous Short Integer Solution* problem relative to the NTRU trapdoor algorithm **TpdGen** with parameters $q, s, B > 0, \alpha \geq 1$ are defined via the games t - $\mathcal{R}$ -ISIS and t - $\mathcal{R}$ -SPISIS, depicted in Figure 5. We define the advantages of $\mathbf{A}$ as

$$\begin{aligned} \text{Adv}_{q,\alpha,B,\mathbf{A}}^{t\text{-}\mathcal{R}\text{-ISIS}} &:= \Pr[t\text{-}\mathcal{R}\text{-ISIS}_{q,\alpha,B}(\mathbf{A}) \Rightarrow 1], \\ \text{Adv}_{q,\alpha,s,B,\mathbf{A}}^{t\text{-}\mathcal{R}\text{-SPISIS}} &:= \Pr[t\text{-}\mathcal{R}\text{-SPISIS}_{q,\alpha,s,B}(\mathbf{A}) \Rightarrow 1]. \end{aligned}$$

According to [LM06], $1\text{-}\mathcal{R}\text{-ISIS}_{q,\alpha,B}$ is as hard as $\mathbf{SVP}_\gamma$ for $\gamma = \tilde{O}(nB)$ over ideal lattices. We define the problems with respect to an NTRU key instead of a uniformly random element, since **ISIS** is not believed to become easier in that case. However, if this should turn out to be wrong, the advantage of our definition can be trivially upper bounded by the sum of the decisional NTRU advantage and the usual ring **ISIS** definition. We do not aim to determine the exact hardness of these somewhat tailored assumptions, and we make the assumption that both t - $\mathcal{R}$ -ISIS and t - $\mathcal{R}$ -SPISIS instances are as hard as random **ISIS** instances, though improved attacks on t - $\mathcal{R}$ -ISIS may exist [Ber22]. Rather, our primary goal is to define assumptions that precisely capture FALCON's security – assumptions that are not merely sufficient but also necessary – thereby providing clear and well-posed targets for cryptanalysts. In particular, any improvements in cryptanalytic attacks against these assumptions directly translates into attacks on the plain or strong unforgeability of FALCON. More precisely, given an attacker against $(Q_H + 1)\text{-}\mathcal{R}\text{-ISIS}$ (Q_H being the number of random oracle queries) one can directly use it to break FALCON's **UF-CMA** security by simply forwarding random oracle queries to the $(Q_H + 1)\text{-}\mathcal{R}\text{-ISIS}$ adversary. The same holds for $Q_s\text{-}\mathcal{R}\text{-SPISIS}$ (Q_s being the number of signing queries) and FALCON's **SUF-CMA** security.

Game $t\text{-}\mathcal{R}\text{-ISIS}_{q,\alpha,B}(\mathbf{A})$	
01	$(\cdot, \cdot, \cdot, \cdot, \mathbf{h}) \xleftarrow{\$} \text{TpdGen}(\alpha, q)$
02	for $i \in [t]$
03	$\mathbf{c}_i \xleftarrow{\$} \mathcal{R}_q$
04	$(j, \mathbf{u}, \mathbf{v}) \xleftarrow{\$} \mathbf{A}(\mathbf{h}, \mathbf{c}_1, \dots, \mathbf{c}_t)$
05	return $\llbracket \mathbf{u} + \mathbf{h} \cdot \mathbf{v} = \mathbf{c}_j \wedge \ (\mathbf{u}, \mathbf{v})\ _2 \leq B \rrbracket$
Game $t\text{-}\mathcal{R}\text{-SPISIS}_{q,\alpha,s,B}(\mathbf{A})$	
06	$(\cdot, \cdot, \cdot, \cdot, \mathbf{h}) \xleftarrow{\$} \text{TpdGen}(\alpha, q)$
07	$\mathbf{\Lambda} := \mathbf{\Lambda}(B_h)$
08	for $i \in [t]$
09	repeat
10	$\mathbf{c}_i \xleftarrow{\$} \mathcal{R}_q$
11	$(\mathbf{u}_i, \mathbf{v}_i) \xleftarrow{\$} \mathcal{D}_{\mathbf{\Lambda}(\mathbf{c}_i, 0), s}$
12	until $\ (\mathbf{u}_i, \mathbf{v}_i)\ _2 \leq B$
13	$(j, \mathbf{u}, \mathbf{v}) \xleftarrow{\$} \mathbf{A}(\mathbf{h}, \{(\mathbf{c}_i, \mathbf{u}_i, \mathbf{v}_i)\}_{i \in [t]})$
14	return $\llbracket \mathbf{u} + \mathbf{h} \cdot \mathbf{v} = \mathbf{c}_j \wedge \ (\mathbf{u}, \mathbf{v})\ _2 \leq B \wedge (\mathbf{u}, \mathbf{v}) \neq (\mathbf{u}_j, \mathbf{v}_j) \rrbracket$

Figure 5. Games defining $t\text{-}\mathcal{R}\text{-ISIS}_{q,\alpha,B}$ and $t\text{-}\mathcal{R}\text{-SPISIS}_{q,\alpha,s,B}$.

3 Security arguments using the Rényi Divergence

We introduce new techniques for applying Rényi arguments to prove the security of FALCON-type schemes. These general results may be useful for a broader class of schemes that rely on the Rényi divergence, with

potential applications to works such as [EFG⁺22, ENS⁺23, GJK24, YJW23]. First, we extend [GPV08, Cor. 2.8], originally stated in terms of statistical distance, to accommodate the Rényi divergence. Such a lemma for Rényi order ∞ was stated in [BLL⁺15, Lem. 2.10]. While these results are not entirely novel, we provide the necessary details for their application in our formal proof. Lemma 8 shows that a Gaussian sample over Λ is distributed almost-uniformly modulo a sublattice Λ' , provided the standard deviation exceeds the smoothing parameter of Λ' .

Lemma 8 (Rényi Divergence of Gaussian Sample over Λ/Λ' (adapted from [GPV08, Cor. 2.8])).

Let Λ, Λ' be n -dimensional full-rank lattices with $\Lambda' \subseteq \Lambda$. Then for any $a \in (1, \infty)$, $\epsilon \in (0, \frac{1}{2})$, any $s \geq \eta_\epsilon(\Lambda')$, and any $c \in \mathbb{R}^n$,

$$R_a(\mathcal{U}(\Lambda/\Lambda') \parallel \mathcal{D}_{\Lambda/\Lambda', s, c}) \lesssim 1 + \frac{2a\epsilon^2}{(1-\epsilon)^2}.$$

The proof can be found in Section B.1. Similarly, we extend [GPV08, Lem. 5.2], also originally stated in terms of statistical distance, to work with the Rényi divergence. The following lemma states that an error vector taken from an appropriate Discrete Gaussian over $\mathbb{Z}^m$ corresponds to a nearly-uniform syndrome.

Lemma 9 (Rényi divergence (adapted from [GPV08, Lem 5.2])). If the columns of $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ generate $\mathbb{Z}_q^n$, $a \in (1, \infty)$, $\epsilon \in (0, \frac{1}{2})$, and $s \geq \eta_\epsilon(\Lambda^\perp(\mathbf{A}))$; then for $e \sim \mathcal{D}_{\mathbb{Z}^m, s}$, the distribution $\mathcal{P} = \mathcal{U}(\mathbb{Z}_q^n)$, and the distribution $\mathcal{Q}$ of the syndromes $u = \mathbf{A}e \bmod q$, it holds that

$$R_a(\mathcal{P} \parallel \mathcal{Q}) \lesssim 1 + \frac{2a\epsilon^2}{(1-\epsilon)^2}.$$

The proof can be found in Section B.2.

Corollary 2 (Rényi uniformity for NTRU). Let q be prime, $\mathbf{h} \in \mathcal{R}_q \setminus \{\mathbf{0}\}$, $a \in (1, \infty)$, $\epsilon \in (0, \frac{1}{2})$, $s \geq \eta_\epsilon(\Lambda_{\mathbf{h}, q})$, $\mathcal{P} = \mathcal{U}(\mathcal{R}_q)$, and $\mathcal{Q}$ the distribution of $\mathbf{u} + \mathbf{v} \cdot \mathbf{h} \bmod q$ where $\mathbf{u}, \mathbf{v} \sim \mathcal{D}_{\mathcal{R}, s}$. Then it holds that

$$R_a(\mathcal{P} \parallel \mathcal{Q}) \lesssim 1 + \frac{2a\epsilon^2}{(1-\epsilon)^2}.$$

The proof can be found in Section B.3.

The next lemma shows that the tailbounds of two distributions with a small relative error are close.¹⁰

Lemma 10 (Relative Error for Tailbounds). Let $\mathcal{P}$ and $\mathcal{Q}$ be two distributions with $\text{sup}(\mathcal{P}) = \text{sup}(\mathcal{Q}) = \mathbb{Z}^n$ and $\delta_{RE}(\mathcal{P}, \mathcal{Q}) = \delta$. Then for any $\beta \geq 0$,

$$\Pr_{x \leftarrow \mathcal{P}}[\|x\|_2 > \beta] \leq \Pr_{x \leftarrow \mathcal{Q}}[\|x\|_2 > \beta] \cdot (1 + \delta).$$

The proof can be found in Section B.4.

For the Rényi divergence, the order a can take any value in $(1, \infty)$, where a smaller a offers better efficiency, and a larger a enables a tighter proof. The description of the lemma is chosen to match statements usually occurring in a security bound (compare for example Section 4.2). For two events $\mathcal{E}_1$ and $\mathcal{E}_2$, Lemma 11 states the minimal number of bits that are lost when moving from $\mathcal{E}_1$ to $\mathcal{E}_2$. Optimising the Rényi order was previously considered in [TT15].

Lemma 11 (Optimal Rényi Order). For $\lambda \in \mathbb{N}$, let $\mathcal{E}_1, \mathcal{E}_2$ be two events such that $\Pr[\mathcal{E}_1] \geq 2^{-\lambda}$. Assume that for any $Q \in \mathbb{N}$, $a \in (1, \infty)$, and $R_a \in [1, \infty)$ it holds that

$$\Pr[\mathcal{E}_2] \leq R_a^Q \cdot \Pr[\mathcal{E}_1]^{\frac{a-1}{a}}.$$

Then

$$-\log(\Pr[\mathcal{E}_2]) \geq -\log(\Pr[\mathcal{E}_1]) - \min_{a > 1} \left\{ Q \log R_a + \frac{\lambda}{a} \right\}.$$

The proof can be found in Section B.5.

¹⁰ Note that the Rényi divergence can be bounded by the relative error using Lemma 6.

4 CoreFalcon⁺: A Framework for Falcon

Let n be a power of 2, q prime, and $\mathcal{R} = \mathbb{Z}[X]/(X^n + 1)$. Let $\alpha \in \mathbb{R}^{>1}$ (basis quality), $\beta \in \mathbb{R}^{>0}$ (signature norm bound), $s \in \mathbb{R}^{>0}$ (Gaussian standard deviation), and $k \in \mathbb{N}$ (size of seed) be fixed parameters. Let $\text{TpGen} : \mathbb{R} \times \mathbb{Z} \rightarrow \mathcal{R}^5$ be a trapdoor generation algorithm, let $\text{PreSmp} : \mathbb{Z}^{2n \times 2n} \times \mathbb{R} \times \mathcal{R}^2 \rightarrow \mathcal{R}^2$ be a preimage sampling algorithm, and $\text{H} : \mathcal{R}_q \times \{0, 1\}^k \times \mathcal{M} \rightarrow \mathcal{R}_q$ be a hash function. The defining algorithms of signature schemes COREFALCON^+ and COREFALCON are given in Figure 6.

<u>Gen</u>	<u>Ver(pk = h, m, σ = (r, s₂))</u>
01 $(f, g, F, G, h) \xleftarrow{\$} \text{TpGen}(\alpha, q)$	11 $c := \text{H}(pk, r, m)$
02 $B := \begin{bmatrix} \mathcal{A}(g) & \mathcal{A}(G) \\ -\mathcal{A}(f) & -\mathcal{A}(F) \end{bmatrix} \in \mathbb{Z}^{2n \times 2n}$	12 $s_1 := c - s_2 \cdot h \pmod q$
03 return $(sk := B, pk := h)$	13 return $\mathbb{I}(\ s_1, s_2\ _2 \leq \beta)$
<u>Sgn(sk = B, m)</u>	<u>Sgn⁺(sk = B, m)</u>
04 $r \xleftarrow{\$} \{0, 1\}^k$	14 repeat
05 $c := \text{H}(pk, r, m) \in \mathcal{R}_q$	15 $r \xleftarrow{\$} \{0, 1\}^k$
06 repeat	16 $c := \text{H}(pk, r, m) \in \mathcal{R}_q$
07 $(s_1, s_2) \xleftarrow{\$} \text{PreSmp}(B, s, (c, 0))$	17 $(s_1, s_2) \xleftarrow{\$} \text{PreSmp}(B, s, (c, 0))$
08 until $\ (s_1, s_2)\ _2 \leq \beta$	18 until $\ (s_1, s_2)\ _2 \leq \beta$
09 $\sigma := (r, s_2) \in \{0, 1\}^k \times \mathcal{R}$	19 $\sigma := (r, s_2) \in \{0, 1\}^k \times \mathcal{R}$
10 return σ	20 return σ

Figure 6. Construction of the $\text{COREFALCON} = (\text{Gen}, \text{Sgn}, \text{Ver})$ and $\text{COREFALCON}^+ = (\text{Gen}, \text{Sgn}^+, \text{Ver})$ signature schemes.

Note that COREFALCON^+ is a slight modification of COREFALCON : In signing Sgn^+ of COREFALCON^+ , picking the random seed r and computing the ring element $c = \text{H}(pk, r, m)$ is performed inside the repeat loop (lines 14–18), while COREFALCON picks a fixed seed r . This modification is not only conceptual; see the discussion below.

The NIST FALCON signature schemes, FALCON-512 and FALCON-1024, can be seen as specific instantiations of COREFALCON .¹¹ Unfortunately, we were not able to analyse the security of COREFALCON since picking the random seed r outside of the repeat loop crucially affects the distribution of the signature in a way we are not able to simulate. Instead, in the next section, we will provide a general security analysis of the COREFALCON^+ framework and derive concrete security levels from modifications FALCON^+ -512 and FALCON^+ -1024.

Note that our modular analysis can be applied to COREFALCON^+ variants that use alternative samplers or key generation procedures, including recent approaches like [EFG⁺22] and [ENS⁺23].

Parameter	Description \ NIST Level	Falcon-512	Falcon-1024
		I	V
n	Degree of ring $\mathcal{R}$	512	1024
q	Modulus	12289	
ϵ	smoothing parameter quality	$2^{-35.5}$	2^{-36}
s	Standard deviation	165.736617183	168.388571447
τ	Tailcut rate	1.1	
β	Max. signature norm bound	5833.93	8382.44
k	Bit size of the salt	320	

Table 2. Parameter sets for FALCON-512/FALCON⁺-512 and FALCON-1024/FALCON⁺-1024 [PFH⁺22, Tab. 3.3].

4.1 Falcon Parameter Sets

As discussed above, FALCON can be seen as COREFALCON with two parameter sets [PFH⁺22]; a smaller set with ring degree $n = 512$ (FALCON-512) targeting NIST security level I, and a larger set with ring degree $n = 1024$ (FALCON-1024), targeting NIST security level V. Both sets use the same modulus $q = 12289$. The smoothing parameter quality is defined as $\epsilon = 1/\sqrt{Q_s \cdot \lambda}$, where Q_s represents the recommend maximum number of signing queries, set to 2^{64} , and λ is the security parameter, set to 128 for NIST level I and 256 for NIST level V. Given ϵ and $\alpha = 1.17$, the standard deviation s is given by

$$s = \frac{1}{\pi} \sqrt{\frac{\ln(4n(1+1/\epsilon))}{2}} \cdot \alpha \sqrt{q}.$$

By Definition 7 and Lemmas 3 and 4, the standard deviation of signatures is lower bounded by the smoothing parameter multiplied by the Gram–Schmidt norm of the trapdoor. The maximum signature norm bound β is set using a fixed tailcut rate $\tau = 1.1$, resulting in $\beta = \tau s \sqrt{2n}$. An overview of the relevant parameters of FALCON-512 and FALCON-1024 can be found in Table 2. We define FALCON⁺-512 and FALCON⁺-1024 using the COREFALCON⁺ framework, instantiated with the parameters from Table 2. FALCON uses the so-called FFO SAMPLER (Fast Fourier Orthogonalization) from [DP16] to instantiate the preimage sampler PreSmp. For completeness we include an analysis of the FFO SAMPLER in Section F.

4.2 Security Bounds for CoreFalcon⁺

In this section, we present two theorems that quantify the concrete security of COREFALCON⁺ in the random oracle model. Theorem 1 provides a security bound for unforgeability. Theorem 2 provides a security bound for *strong* unforgeability but relies on a stronger assumption.

Theorem 1 (Unforgeability). For any adversary A against the **UF-CMA** security of COREFALCON⁺ (Figure 6) running in time t_A , making at most Q_s signing queries and Q_H random oracle queries, there exists an adversary B against $(Q_H + 1)$ - $\mathcal{R}$ -ISIS running in time $t_B \approx t_A$ such that for all $C_s \in \mathbb{N}^{\geq 1}$ and $a_u, a_p \in \mathbb{R}^{>1}$

¹¹ In the signing process for FALCON-512 and FALCON-1024, a (public) compression technique is applied to the signature, and the loop is repeated until the signature reaches the desired compression level. This modification is mainly conceptual, as with the parameters of FALCON, the compressed signature typically reaches a sufficiently small size with high probability. Furthermore, COREFALCON includes the public key in the hash function H , whereas FALCON-512 and FALCON-1024 do not. Including the public key in the hash function to make it key-contributory is generally considered good cryptographic engineering. Moreover, including the public key in the hash, as in the Pornin-Stern transformation [PS05], has been shown to provide additional security properties beyond unforgeability [CDF⁺21, DFF24].

it holds

$$\begin{aligned} \text{Adv}_{\text{COREFALCON}^+, \mathbf{A}}^{Q_s\text{-}\mathbf{UF}\text{-}\mathbf{CMA}} &\leq \left(r_u^{C_s} \cdot \left(r_p^{C_s} \cdot \text{Adv}_{q, \alpha, \beta, \mathbf{B}}^{(Q_H+1)\text{-}\mathcal{R}\text{-}\mathbf{ISIS}} \right)^{\frac{a_p-1}{a_p}} \right)^{\frac{a_u-1}{a_u}} \\ &\quad + \sum_{i=0}^{Q_s} \binom{C_s}{i} (1 - p_{\text{PreSmp}, \beta})^{C_s-i} (p_{\text{PreSmp}, \beta})^i + \frac{C_s(Q_H + C_s)}{2^k}, \end{aligned}$$

where

$$\begin{aligned} p_{\text{PreSmp}, \beta} &:= \min_{(\mathbf{B}, \cdot) \in \text{TpdGen}, \mathbf{c} \in \mathcal{R}_q} \Pr_{(\mathbf{s}_1, \mathbf{s}_2) \leftarrow \mathbb{S}\text{PreSmp}(\mathbf{B}, s, (\mathbf{c}, \mathbf{0}))} [\|\mathbf{s}_1, \mathbf{s}_2\|_2 \leq \beta], \\ r_u &= \max_{(\cdot, \mathbf{h}) \in \text{TpdGen}} R_{a_u}(\mathcal{P} \parallel \mathcal{Q}_{\mathbf{h}}) \text{ with } \mathcal{P} = \mathcal{U}(\mathcal{R}_q) \text{ and } \mathcal{Q}_{\mathbf{h}} \text{ the distribution of } \mathbf{s}_1 + \mathbf{s}_2 \cdot \mathbf{h} \text{ mod } q, \text{ where} \\ &\quad \mathbf{s}_1, \mathbf{s}_2 \sim \mathcal{D}_{\mathcal{R}, s}, \\ r_p &= \max_{(\mathbf{B}, \cdot) \in \text{TpdGen}, \mathbf{c} \in \mathcal{R}_q} R_{a_p}(\text{PreSmp}(\mathbf{B}, s, (\mathbf{c}, \mathbf{0})) \parallel \mathcal{D}_{\Lambda, s}) \text{ with } \Lambda = \Lambda(\mathbf{B})_{(\mathbf{c}, \mathbf{0})}. \end{aligned}$$

REMARK. Note that the bound of Theorem 1 (and Theorem 2) holds for all choices of constants $C_s \in \mathbb{N}^{\geq 1}$ and $a_u, a_p \in \mathbb{R}^{>1}$. We will refer to these as *proof constants*. In Section 5, we will derive optimal choices for these proof constants that minimise the security loss for concrete and relevant instantiations of COREFALCON^+ . The proof of Theorem 1 can be found in Section C.

Interestingly, the hardness of $(Q_H+1)\text{-}\mathcal{R}\text{-}\mathbf{ISIS}$ is not only sufficient for Theorem 1, but it is also necessary. Specifically, an attack on $(Q_H+1)\text{-}\mathcal{R}\text{-}\mathbf{ISIS}$ would directly lead to an attack on the $\mathbf{UF}\text{-}\mathbf{CMA}$ security of FALCON . Similarly, Theorem 2 requires the hardness of $Q_s\text{-}\mathcal{R}\text{-}\mathbf{SPISIS}$, and an attack on this would result in an attack on the $\mathbf{SUF}\text{-}\mathbf{CMA}$ security of FALCON .

Theorem 2 (Strong Unforgeability). For any adversary $\mathbf{A}$ against the $\mathbf{SUF}\text{-}\mathbf{CMA}$ security of COREFALCON^+ (Figure 6) running in time $t_{\mathbf{A}}$, making at most Q_s signing queries and Q_H random oracle queries, there exist an adversary $\mathbf{B}$ against $Q_s\text{-}\mathcal{R}\text{-}\mathbf{SPISIS}$ running in time $t_{\mathbf{B}} \approx t_{\mathbf{A}}$ such that for all $C_s \in \mathbb{N}^{\geq 1}$ and $a_p \in \mathbb{R}^{>1}$ it holds

$$\begin{aligned} \text{Adv}_{\text{COREFALCON}^+, \mathbf{A}}^{Q_s\text{-}\mathbf{SUF}\text{-}\mathbf{CMA}} &\leq \text{Adv}_{\text{COREFALCON}^+, \mathbf{A}}^{Q_s\text{-}\mathbf{UF}\text{-}\mathbf{CMA}} + \left(r_p^{C_s} \cdot \left(\text{Adv}_{q, \alpha, s, \beta, \mathbf{B}}^{Q_s\text{-}\mathcal{R}\text{-}\mathbf{SPISIS}} + p_{\text{binom}} \right) \right)^{\frac{a_p-1}{a_p}} \\ &\quad + p_{\text{binom}} + \left(\frac{Q_s + 1}{2p_{\text{PreSmp}, \beta}^2} + \frac{2Q_H}{p_{\text{PreSmp}, \beta}} \right) Q_s 2^{-k}, \end{aligned}$$

where

$$\begin{aligned} p_{\text{PreSmp}, \beta} &:= \min_{(\mathbf{B}, \cdot) \in \text{TpdGen}, \mathbf{c} \in \mathcal{R}_q} \Pr_{(\mathbf{s}_1, \mathbf{s}_2) \leftarrow \mathbb{S}\text{PreSmp}(\mathbf{B}, s, (\mathbf{c}, \mathbf{0}))} [\|\mathbf{s}_1, \mathbf{s}_2\|_2 \leq \beta], \\ p_{\text{binom}} &:= \sum_{i=0}^{Q_s} \binom{C_s}{i} (1 - p_{\text{PreSmp}, \beta})^{C_s-i} (p_{\text{PreSmp}, \beta})^i \\ r_p &= \max_{(\mathbf{B}, \cdot) \in \text{TpdGen}, \mathbf{c} \in \mathcal{R}_q} R_{a_p}(\text{PreSmp}(\mathbf{B}, s, (\mathbf{c}, \mathbf{0})) \parallel \mathcal{D}_{\Lambda, s}) \text{ with } \Lambda = \Lambda(\mathbf{B})_{(\mathbf{c}, \mathbf{0})}. \end{aligned}$$

The proof of Theorem 2 can be found in Section D.

5 Parameters and Analysing the Security Bound

In this section, we analyse the concrete security bounds for $\text{FALCON}^+\text{-512}$ and $\text{FALCON}^+\text{-1024}$ from Section 4.1. Recall that $\text{FALCON}^+\text{-512}$ and $\text{FALCON}^+\text{-1024}$ are slight modifications of $\text{FALCON}\text{-512}$ and $\text{FALCON}\text{-1024}$, respectively (with the same parameter sets), where signing includes picking the random seed inside of the repeat loop. Concretely, we will use the Theorems from Section 4.2 to derive the proof constants C_s , a_u , and a_p for an optimal tightness of the security proofs. The FALCON specification suggests setting the Rényi order to $a_p = 2\lambda$, which is sufficient, but not ideal. We proceed as follows: first, we estimate the $t\text{-}\mathcal{R}\text{-}\mathbf{ISIS}/t\text{-}\mathcal{R}\text{-}\mathbf{SPISIS}$ bit security. Next, we analyse the bound in Theorem 1 and Theorem 2, beginning with proof constant C_s , denoting the maximal repetitions in the signing oracle. Next, based on the bit security of the $t\text{-}\mathcal{R}\text{-}\mathbf{ISIS}/t\text{-}\mathcal{R}\text{-}\mathbf{SPISIS}$ term, we iteratively apply the Rényi arguments, carefully choosing the optimal orders a_u and a_p to minimise the security loss. Finally, we combine all results to calculate the final bit security, presenting an overview in Table 3, followed by a discussion of the findings.

5.1 Security of $t\mathcal{R}$ -ISIS and $t\mathcal{R}$ -SPISIS

We estimate the security of the $t\mathcal{R}$ -ISIS and $t\mathcal{R}$ -SPISIS terms in our bounds. We consider the $t\mathcal{R}$ -ISIS and $t\mathcal{R}$ -SPISIS problems (as defined in Definition 10), parametrised by a trapdoor generation algorithm TpdGen with trapdoor quality α and modulus q . For plain unforgeability, Theorem 1 provides a reduction to $t\mathcal{R}$ -ISIS with a norm bound of β . For strong unforgeability, Theorem 2 gives a reduction to $t\mathcal{R}$ -SPISIS with the same norm bound. For the hardness of $t\mathcal{R}$ -ISIS / $t\mathcal{R}$ -SPISIS we use a ring dimension of $n = 512$ ($n = 1024$) and modulus $q = 12289$. The length bound $\beta = \tau s \sqrt{2n}$ results in $\beta_I = 5833.93$ for FALCON⁺-512 and $\beta_V = 8382.44$ for FALCON⁺-1024 (see Table 2). We make the assumption that $t\mathcal{R}$ -ISIS and $t\mathcal{R}$ -SPISIS instances are as hard as random SIS instances. Although it is possible that there are more efficient attacks against $t\mathcal{R}$ -ISIS [Ber22], we argue that a direct reduction to $t\mathcal{R}$ -ISIS in Theorem 1 is meaningful, as it most accurately captures the security of the scheme. That is, $t\mathcal{R}$ -ISIS does not only suffice for *plain unforgeability*, but is, in fact, also necessary. Specifically, an attack on $t\mathcal{R}$ -ISIS would directly imply an attack on the *plain unforgeability* of FALCON. The same applies to $t\mathcal{R}$ -SPISIS and the *strong unforgeability* of FALCON. We estimate the security of SIS using the “lattice-estimator” [APS15a, APS15b] with the `SIS.estimate.rough()` function, which computes the concrete bit security based on the *core-SVP methodology* from [ADPS16].¹² The resulting levels of bit security are summarised in Table 3. We refer to Figure 11 in Section E for the concrete prompts of the lattice estimator.

5.2 Number Of Signing Repetitions C_s

The proof constant C_s defines the maximum number of repetitions to the signing oracle. Increasing C_s inflates all terms in the security bound, except for the binomial term. Hence, to obtain an optimal bound that fulfils the target security level λ , we have to find the smallest C_s such that the binomial term is less than $2^{-\lambda}$. The following lemma establishes this for FALCON⁺-512 and FALCON⁺-1024.

Lemma 12 (Optimal C_s). For FALCON⁺-512 with $\lambda = 128$ it holds that,

$$\arg \min_{C_s} \left\{ C_s \left| \sum_{i=0}^{Q_s} \binom{C_s}{i} (1 - p_{\text{PreSmp}, \beta})^{C_s-i} (p_{\text{PreSmp}, \beta})^i \leq 2^\lambda \right. \right\} \lesssim 2^{64} + 2^{50},$$

and for FALCON⁺-1024 with $\lambda = 256$ it holds that

$$\arg \min_{C_s} \left\{ C_s \left| \sum_{i=0}^{Q_s} \binom{C_s}{i} (1 - p_{\text{PreSmp}, \beta})^{C_s-i} (p_{\text{PreSmp}, \beta})^i \leq 2^\lambda \right. \right\} \lesssim 2^{64} + 2^{36}.$$

The proof can be found in Section E.1. For different values Q_s , C_s can be computed in the same way, as shown in Table 3.

5.3 Rényi Terms

FALCON builds on the work of [Pre17, Lem. 6] which suggests that setting $a_p = 2\lambda$ “seems to be good compromise”. Although this is true for certain problem instantiations, Lemma 11 makes this choice less ad hoc and allows us to set the order of the Rényi divergence optimally, similar to [TT15]. We start with optimising the Rényi order for the unforgeability bound (Theorem 1), i.e., the reduction to $t\mathcal{R}$ -ISIS.

FALCON⁺-512. We start with the advantage for $t\mathcal{R}$ -ISIS which gives 120 bits security, so for the inner most part of the bound we have to preserve at most $\lambda = 120$ bits of security.

¹² We acknowledge that other tools for estimating the hardness of lattice problems exist [DDGR20, Duc20], and work has been done to analyse the hardness of ISIS for small moduli [DEP23]. Any improvements in the cryptanalysis of the underlying problems would also lead to improved attacks on the scheme, in which case our theorem bounds would remain unchanged, and only Table 3 would need to be updated.

Corollary 3 (Rényi Loss for Falcon⁺-512 (Preimage Sampler) in Thm. 1). For $\varepsilon \geq 2^{-\lambda} = 2^{-120}$, $r_p = R_{a_p}(\text{PreSmp} \parallel \mathcal{D})$, $C_s = 2^{64} + 2^{50}$, and the parameters for FALCON⁺-512, the Rényi argument for

$$r_p^{C_s} \varepsilon^{\frac{a_p-1}{a_p}}$$

loses at most 3.5 bits for an order $a_p \approx 72.96$.

The proof can be found in Section E.2. Next, we consider the 3.5 bits lost from Corollary 3 when analysing the bits lost for the uniformity result.

Corollary 4 (Rényi Loss for Falcon⁺-512 (Uniformity) in Thm. 1). For $\varepsilon \geq 2^{-\lambda} = 2^{-116.5}$, $r_u = R_{a_u}(\mathcal{U}(\mathcal{R}_q) \parallel \mathcal{U}_h)$, $C_s = 2^{64} + 2^{50}$, and the parameters for FALCON⁺-512, the Rényi argument for

$$r_u^{C_s} \varepsilon^{\frac{a_u-1}{a_u}}$$

loses at most 3.5 bits for an order $a_u \approx 71.73$.

The proof can be found in Section E.3.

FALCON⁺-1024. We apply the same arguments as for FALCON⁺-512. The analogous corollaries can be found in Section E.4.

OTHER BOUNDS AND NUMBER OF SIGNING QUERIES. The optimal Rényi orders for the strong unforgeability bound (Theorem 2) as well as for different choices of the maximum number of signing queries Q_s can be computed in the same way. We give an overview in the following section.

5.4 Final Security and Discussion

Corollaries 3 to 6 show that $C_s < 2Q_s$. To conclude the analysis of the bounds, we note that the term $C_s(C_s + Q_H)/2^k$ provides λ bits of security when $k \geq \log(2Q_s) + \lambda$. For both parameter sets, FALCON⁺ achieves this by its choice of $k = 320$ up to a loss of less than a bit. The binomial term fulfils λ bits of security by choosing an appropriate C_s , as detailed in the proof of Lemma 12. An overview of the results from the previous subsections is presented in Table 3 for FALCON⁺-512 and FALCON⁺-1024. Note that while the computational term in the bound for FALCON⁺-1024 ensures 270 bits of security, the statistical terms described above limit the overall security to 256 bits. Below, we address key findings and issues, suggesting possible solutions.

STRONG UNFORGEABILITY. We assume that both $t\mathcal{R}$ -SPISIS (and $t\mathcal{R}$ -ISIS) are as hard as plain SIS. Comparing the security bounds from Theorem 1 and Theorem 2, one can observe that the dominating term in Theorem 2 is the **UF-CMA** term. Therefore, the provable bit security levels for the strong unforgeability of FALCON⁺-512 and FALCON⁺-1024 are essentially the same as in Table 3.

NUMBER OF SIGNING QUERIES. For FALCON⁺-512, we provide bit security estimates for both reduced and full 2^{64} signing queries, as required by NIST. Allowing 2^{64} queries increases the Rényi divergence loss, which is problematic, since the security of SIS is already tightly set to the target level. In contrast, FALCON⁺-1024 benefits from a larger security margin between SIS and the target security, making it more tolerant to larger Rényi losses. Therefore, we also present the maximum number of signing queries that can be supported while maintaining a Rényi loss of at most 1 bit. While increasing the smoothing parameter error ϵ could help maintain tight Rényi bounds even up to 2^{64} queries, doing so would increase other parameters – such as the signature size.

Acknowledgements. The authors thank Brice Minaud for his suggestion of using second-preimage resistance for strong unforgeability. They also thank Doreen Riepel, Melissa Chase, Mihir Bellare, Joel Gärtner, Thomas Prest, and Damien Stehlé for helpful feedback and comments. Part of the work was done while Phillip Gajland was affiliated with RUB and MPI-SP. Phillip Gajland was supported by the

Table 3. Provable security levels of FALCON^+ -512 and FALCON^+ -1024. The $\star$ symbol at 270 bits refers to the bit security of the computational term. For further details, see Section 5.4.

Parameters	UF-CMA (Thm. 1)		
	FALCON^+ -512		FALCON^+ -1024
Ring $\mathcal{R}_q$	$\mathbb{Z}_{12289}[X]/(X^{512} + 1)$		$\mathbb{Z}_{12289}[X]/(X^{1024} + 1)$
t - $\mathcal{R}$ -ISIS length bound β	5833.93		8382.44
Bit security (core-SVP) t - $\mathcal{R}$ -ISIS $_{q=q, \alpha=1.17, B=\beta}$	120		278
Max Signing queries Q_s	2^{58}	2^{64}	2^{64}
Max repetitions $C_s(\lambda, Q_s)$	$2^{58} + 2^{44}$	$2^{64} + 2^{50}$	$2^{64} + 2^{36}$
Rényi Order a_p	583.67	72.96	157.05
Rényi Order a_u	582.46	71.73	155.92
Bits lost from Rényi a_p	0.5	3.5	4
Bits lost from Rényi a_u	0.5	3.5	4
Final bit security	119	113	256 (270)$\star$

Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC 2092 CASA - 390781972. Jonas Janneck was supported by the European Union (ERC AdG REWORC - 101054911). Eike Kiltz was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy – EXC 2092 CASA - 390781972, and by the European Union (ERC AdG REWORC - 101054911).

References

- [ADPS16] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange - A new hope. In Thorsten Holz and Stefan Savage, editors, *USENIX Security 2016: 25th USENIX Security Symposium*, pages 327–343, Austin, TX, USA, August 10–12, 2016. USENIX Association. URL: <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/alkim>. (Cited on page 19.)
- [ADRS15] Divesh Aggarwal, Daniel Dadush, Oded Regev, and Noah Stephens-Davidowitz. Solving the shortest vector problem in 2^n time using discrete Gaussian sampling: Extended abstract. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *47th Annual ACM Symposium on Theory of Computing*, pages 733–742, Portland, OR, USA, June 14–17, 2015. ACM Press. doi:10.1145/2746539.2746606. (Cited on page 11.)
- [Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *28th Annual ACM Symposium on Theory of Computing*, pages 99–108, Philadelphia, PA, USA, May 22–24, 1996. ACM Press. doi:10.1145/237814.237838. (Cited on pages 3 and 5.)
- [APS15a] Martin R. Albrecht, Rachel Player, and Sam Scott. Lattice estimator. <https://github.com/malb/lattice-estimator>, 2015. Commit: 14a362513c9197dd959bc72428425abe0309779a. (Cited on pages 8 and 19.)
- [APS15b] Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of learning with errors. *Journal of Mathematical Cryptology*, 9(3):169–203, 2015. URL: <https://doi.org/10.1515/jmc-2015-0016> [cited 2024-05-23], doi:doi:10.1515/jmc-2015-0016. (Cited on pages 8 and 19.)
- [Ban93] W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(1):625–635, December 1993. doi:10.1007/bf01445125. (Cited on page 11.)
- [BBD⁺23] Manuel Barbosa, Gilles Barthe, Christian Doczkal, Jelle Don, Serge Fehr, Benjamin Grégoire, Yu-Hsuan Huang, Andreas Hülsing, Yi Lee, and Xiaodi Wu. Fixing and mechanizing the security proof of Fiat-Shamir with aborts and Dilithium. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023, Part V*, volume 14085 of *Lecture Notes in Computer Science*, pages 358–389, Santa Barbara, CA, USA, August 20–24, 2023. Springer, Cham, Switzerland. doi:10.1007/978-3-031-38554-4_12. (Cited on pages 8 and 33.)

- [BDF⁺11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In Dong Hoon Lee and Xiaoyun Wang, editors, *Advances in Cryptology – ASIACRYPT 2011*, volume 7073 of *Lecture Notes in Computer Science*, pages 41–69, Seoul, South Korea, December 4–8, 2011. Springer Berlin Heidelberg, Germany. doi:10.1007/978-3-642-25385-0_3. (Cited on page 5.)
- [Ber22] Daniel J. Bernstein. Multi-ciphertext security degradation for lattices. Cryptology ePrint Archive, Report 2022/1580, 2022. URL: <https://eprint.iacr.org/2022/1580>. (Cited on pages 14 and 19.)
- [Beu22] Ward Beullens. Breaking rainbow takes a weekend on a laptop. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part II*, volume 13508 of *Lecture Notes in Computer Science*, pages 464–479, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Cham, Switzerland. doi:10.1007/978-3-031-15979-4_16. (Cited on page 3.)
- [BLL⁺15] Shi Bai, Adeline Langlois, Tancrede 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, *Advances in Cryptology – ASIACRYPT 2015, Part I*, volume 9452 of *Lecture Notes in Computer Science*, pages 3–24, Auckland, New Zealand, November 30 – December 3, 2015. Springer Berlin Heidelberg, Germany. doi:10.1007/978-3-662-48797-6_1. (Cited on pages 3, 12, and 15.)
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93: 1st Conference on Computer and Communications Security*, pages 62–73, Fairfax, Virginia, USA, November 3–5, 1993. ACM Press. doi:10.1145/168588.168596. (Cited on pages 5 and 9.)
- [BR96] Mihir Bellare and Phillip Rogaway. The exact security of digital signatures: How to sign with RSA and Rabin. In Ueli M. Maurer, editor, *Advances in Cryptology – EUROCRYPT’96*, volume 1070 of *Lecture Notes in Computer Science*, pages 399–416, Saragossa, Spain, May 12–16, 1996. Springer Berlin Heidelberg, Germany. doi:10.1007/3-540-68339-9_34. (Cited on pages 3 and 5.)
- [BR06] Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In Serge Vaudenay, editor, *Advances in Cryptology – EUROCRYPT 2006*, volume 4004 of *Lecture Notes in Computer Science*, pages 409–426, St. Petersburg, Russia, May 28 – June 1, 2006. Springer Berlin Heidelberg, Germany. doi:10.1007/11761679_25. (Cited on page 9.)
- [CD23] Wouter Castryck and Thomas Decru. An efficient key recovery attack on SIDH. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023, Part V*, volume 14008 of *Lecture Notes in Computer Science*, pages 423–447, Lyon, France, April 23–27, 2023. Springer, Cham, Switzerland. doi:10.1007/978-3-031-30589-4_15. (Cited on page 3.)
- [CDF⁺21] Cas Cremers, Samed Düzlü, Rune Fiedler, Marc Fischlin, and Christian Janson. BUFFing signature schemes beyond unforgeability and the case of post-quantum signatures. In *2021 IEEE Symposium on Security and Privacy*, pages 1696–1714, San Francisco, CA, USA, May 24–27, 2021. IEEE Computer Society Press. doi:10.1109/SP40001.2021.00093. (Cited on page 17.)
- [DD18] Daniel Dadush and Léo Ducas. Determinants, packing and covering, and the minkowski theorems, 2018. URL: <https://homepages.cwi.nl/~dadush/teaching/lattices-2018/notes/lecture-2.pdf>. (Cited on page 27.)
- [DDGR20] Dana Dachman-Soled, Léo Ducas, Huijing Gong, and Mélissa Rossi. LWE with side information: Attacks and concrete security estimation. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology – CRYPTO 2020, Part II*, volume 12171 of *Lecture Notes in Computer Science*, pages 329–358, Santa Barbara, CA, USA, August 17–21, 2020. Springer, Cham, Switzerland. doi:10.1007/978-3-030-56880-1_12. (Cited on page 19.)
- [DEP23] Léo Ducas, Thomas Espitau, and Eamonn W. Postlethwaite. Finding short integer solutions when the modulus is small. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023, Part III*, volume 14083 of *Lecture Notes in Computer Science*, pages 150–176, Santa Barbara, CA, USA, August 20–24, 2023. Springer, Cham, Switzerland. doi:10.1007/978-3-031-38548-3_6. (Cited on page 19.)
- [DFF24] Samed Düzlü, Rune Fiedler, and Marc Fischlin. BUFFing FALCON without increasing the signature size. Cryptology ePrint Archive, Report 2024/710, 2024. URL: <https://eprint.iacr.org/2024/710>. (Cited on page 17.)
- [DLP14] Léo Ducas, Vadim Lyubashevsky, and Thomas Prest. Efficient identity-based encryption over NTRU lattices. In Palash Sarkar and Tetsu Iwata, editors, *Advances in Cryptology – ASIACRYPT 2014, Part II*, volume 8874 of *Lecture Notes in Computer Science*, pages 22–41, Kaoshiung, Taiwan, R.O.C.,

- December 7–11, 2014. Springer Berlin Heidelberg, Germany. doi:10.1007/978-3-662-45608-8_2. (Cited on pages 3, 6, and 11.)
- [DP16] Léo Ducas and Thomas Prest. Fast fourier orthogonalization. In *Proceedings of the 2016 ACM International Symposium on Symbolic and Algebraic Computation*, ISSAC '16, page 191–198, New York, NY, USA, 2016. Association for Computing Machinery. doi:10.1145/2930889.2930923. (Cited on pages 3, 4, 13, 17, 41, and 42.)
- [DRSD14] Daniel Dadush, Oded Regev, and Noah Stephens-Davidowitz. On the closest vector problem with a distance guarantee. In *Proceedings of the 2014 IEEE 29th Conference on Computational Complexity*, CCC '14, page 98–109, USA, 2014. IEEE Computer Society. doi:10.1109/CCC.2014.18. (Cited on page 11.)
- [Duc20] Léo Ducas. Leaky lwe estimator. <https://github.com/lducas/leaky-LWE-Estimator>, 2020. Commit: 0a9caf8bf0f80097724e0c6147194c52c6b90f86. (Cited on page 19.)
- [EFG⁺22] Thomas Espitau, Pierre-Alain Fouque, François Gérard, Mélissa Rossi, Akira Takahashi, Mehdi Tibouchi, Alexandre Wallet, and Yang Yu. Mitaka: A simpler, parallelizable, maskable variant of falcon. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology – EUROCRYPT 2022, Part III*, volume 13277 of *Lecture Notes in Computer Science*, pages 222–253, Trondheim, Norway, May 30 – June 3, 2022. Springer, Cham, Switzerland. doi:10.1007/978-3-031-07082-2_9. (Cited on pages 3, 15, and 16.)
- [ENS⁺23] Thomas Espitau, Thi Thu Quyen Nguyen, Chao Sun, Mehdi Tibouchi, and Alexandre Wallet. Antrag: Annular NTRU trapdoor generation - making mitaka as secure as falcon. In Jian Guo and Ron Steinfeld, editors, *Advances in Cryptology – ASIACRYPT 2023, Part VII*, volume 14444 of *Lecture Notes in Computer Science*, pages 3–36, Guangzhou, China, December 4–8, 2023. Springer, Singapore, Singapore. doi:10.1007/978-981-99-8739-9_1. (Cited on pages 3, 15, and 16.)
- [ENST23] Thomas Espitau, Guilhem Niot, Chao Sun, and Mehdi Tibouchi. SQUIRRELS — Square Unstructured Integer Euclidean Lattice Signature. Technical report, National Institute of Standards and Technology, 2023. available at <https://csrc.nist.gov/Projects/pqc-dig-sig/round-1-additional-signatures>. (Cited on page 4.)
- [FFH25] Pouria Fallahpour, Serge Fehr, and Yu-Hsuan Huang. Tighter quantum security for fiat-shamir-with-aborts and hash-and-sign-with-retry signatures. Cryptology ePrint Archive, Report 2025/985, 2025. URL: <https://eprint.iacr.org/2025/985>. (Cited on page 8.)
- [GJK24] Phillip Gajland, Jonas Janneck, and Eike Kiltz. Ring signatures for deniable AKEM: Gandalf’s fellowship. In Leonid Reyzin and Douglas Stebila, editors, *Advances in Cryptology – CRYPTO 2024, Part I*, volume 14920 of *Lecture Notes in Computer Science*, pages 305–338, Santa Barbara, CA, USA, August 18–22, 2024. Springer, Cham, Switzerland. doi:10.1007/978-3-031-68376-3_10. (Cited on pages 3 and 15.)
- [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 Annual ACM Symposium on Theory of Computing*, pages 197–206, Victoria, BC, Canada, May 17–20, 2008. ACM Press. doi:10.1145/1374376.1374407. (Cited on pages 3, 5, 7, 11, 13, 15, 26, 27, 39, and 40.)
- [HBD⁺22] Andreas Hülsing, Daniel J. Bernstein, Christoph Dobraunig, Maria Eichlseder, Scott Fluhrer, Stefan-Lukas Gazdag, Panos Kampanakis, Stefan Kölbl, Tanja Lange, Martin M. Lauridsen, Florian Mendel, Ruben Niederhagen, Christian Rechberger, Joost Rijneveld, Peter Schwabe, Jean-Philippe Aumasson, Bas Westerbaan, and Ward Beullens. SPHINCS⁺. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>. (Cited on page 3.)
- [HHP⁺03] Jeffrey Hoffstein, Nick Howgrave-Graham, Jill Pipher, Joseph H. Silverman, and William Whyte. NTRUSIGN: Digital signatures using the NTRU lattice. In Marc Joye, editor, *Topics in Cryptology – CT-RSA 2003*, volume 2612 of *Lecture Notes in Computer Science*, pages 122–140, San Francisco, CA, USA, April 13–17, 2003. Springer Berlin Heidelberg, Germany. doi:10.1007/3-540-36563-X_9. (Cited on pages 3 and 6.)
- [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. (Cited on pages 3, 5, and 6.)
- [HPRR20] James Howe, Thomas Prest, Thomas Ricosset, and Mélissa Rossi. Isochronous gaussian sampling: From inception to implementation. In Jintai Ding and Jean-Pierre Tillich, editors, *Post-Quantum Cryptography - 11th International Conference, PQCrypto 2020*, pages 53–71, Paris, France, April 15–17, 2020. Springer, Cham, Switzerland. doi:10.1007/978-3-030-44223-1_4. (Cited on page 26.)

- [HPS98] Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In *Third Algorithmic Number Theory Symposium (ANTS)*, volume 1423 of *Lecture Notes in Computer Science*, pages 267–288. Springer, June 1998. (Cited on pages 3, 6, and 11.)
- [Kim16] Kevin Kimball. Announcing request for nominations for public-key post-quantum cryptographic algorithms. Technical report, National Institute of Standards and Technology, 2016. available at <https://www.federalregister.gov/d/2016-30615>. (Cited on page 3.)
- [Kle00] Philip N. Klein. Finding the closest lattice vector when it’s unusually close. In David B. Shmoys, editor, *11th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 937–941, San Francisco, CA, USA, January 9–11, 2000. ACM-SIAM. (Cited on pages 6, 13, 38, and 39.)
- [LAZ19] Xingye Lu, Man Ho Au, and Zhenfei Zhang. Raptor: A practical lattice-based (linkable) ring signature. In Robert H. Deng, Valérie Gauthier-Umaña, Martín Ochoa, and Moti Yung, editors, *ACNS 19: 17th International Conference on Applied Cryptography and Network Security*, volume 11464 of *Lecture Notes in Computer Science*, pages 110–130, Bogota, Colombia, June 5–7, 2019. Springer, Cham, Switzerland. doi:10.1007/978-3-030-21568-2_6. (Cited on page 3.)
- [LDK⁺22] Vadim Lyubashevsky, Léo Ducas, Eike Kiltz, Tancrede 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>. (Cited on page 3.)
- [LM06] Vadim Lyubashevsky and Daniele Micciancio. Generalized compact Knapsacks are collision resistant. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *ICALP 2006: 33rd International Colloquium on Automata, Languages and Programming, Part II*, volume 4052 of *Lecture Notes in Computer Science*, pages 144–155, Venice, Italy, July 10–14, 2006. Springer Berlin Heidelberg, Germany. doi:10.1007/11787006_13. (Cited on page 14.)
- [LPR13] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for ring-LWE cryptography. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology – EUROCRYPT 2013*, volume 7881 of *Lecture Notes in Computer Science*, pages 35–54, Athens, Greece, May 26–30, 2013. Springer Berlin Heidelberg, Germany. doi:10.1007/978-3-642-38348-9_3. (Cited on pages 3 and 6.)
- [LSS14] Adeline Langlois, Damien Stehlé, and Ron Steinfeld. GGHLite: More efficient multilinear maps from ideal lattices. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology – EUROCRYPT 2014*, volume 8441 of *Lecture Notes in Computer Science*, pages 239–256, Copenhagen, Denmark, May 11–15, 2014. Springer Berlin Heidelberg, Germany. doi:10.1007/978-3-642-55220-5_14. (Cited on page 26.)
- [Lyu12] Vadim Lyubashevsky. Lattice signatures without trapdoors. In David Pointcheval and Thomas Johansson, editors, *Advances in Cryptology – EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 738–755, Cambridge, UK, April 15–19, 2012. Springer Berlin Heidelberg, Germany. doi:10.1007/978-3-642-29011-4_43. (Cited on page 11.)
- [MMP⁺23] Luciano Maino, Chloe Martindale, Lorenz Panny, Giacomo Pope, and Benjamin Wesolowski. A direct key recovery attack on SIDH. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023, Part V*, volume 14008 of *Lecture Notes in Computer Science*, pages 448–471, Lyon, France, April 23–27, 2023. Springer, Cham, Switzerland. doi:10.1007/978-3-031-30589-4_16. (Cited on page 3.)
- [MR04] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. In *45th Annual Symposium on Foundations of Computer Science*, pages 372–381, Rome, Italy, October 17–19, 2004. IEEE Computer Society Press. doi:10.1109/FOCS.2004.72. (Cited on pages 12 and 41.)
- [MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM Journal on Computing*, 37(1):267–302, 2007. arXiv:<https://doi.org/10.1137/S0097539705447360>, doi:10.1137/S0097539705447360. (Cited on pages 11 and 12.)
- [PFH⁺20] Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. FALCON. Technical report, National Institute of Standards and Technology, 2020. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/post-quantum-cryptography-standardization/round-3-submissions>. (Cited on page 3.)
- [PFH⁺22] Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. FALCON. Technical report, National Institute of Standards and Technology, 2022. available at <https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022>. (Cited on pages 6 and 17.)

- [Por25a] Thomas Pornin. Falcon on ARM cortex-m4: an update. Cryptology ePrint Archive, Paper 2025/123, 2025. URL: <https://eprint.iacr.org/2025/123>. (Cited on pages 4 and 6.)
- [Por25b] Thomas Pornin. Fn-dsa (in c). <https://github.com/pornin/c-fn-dsa>, 2025. Commit: 1cdc9c5bdd5b5894475febd7e23abbbcb5056197b. (Cited on pages 4 and 6.)
- [Por25c] Thomas Pornin. Pqc-forum: Constant time digital signature algorithm. Google Groups, Feb 2025. URL: <https://groups.google.com/a/list.nist.gov/g/pqc-forum/c/i38i7t950Cw/m/TanXQ30xAQAJ>. (Cited on page 6.)
- [PP19] Thomas Pornin and Thomas Prest. More efficient algorithms for the NTRU key generation using the field norm. In Dongdai Lin and Kazue Sako, editors, *PKC 2019: 22nd International Conference on Theory and Practice of Public Key Cryptography, Part II*, volume 11443 of *Lecture Notes in Computer Science*, pages 504–533, Beijing, China, April 14–17, 2019. Springer, Cham, Switzerland. doi:10.1007/978-3-030-17259-6_17. (Cited on page 3.)
- [Pre15] Thomas Prest. *Gaussian sampling in lattice-based cryptography*. PhD thesis, Ecole normale supérieure-ENS PARIS, 2015. (Cited on page 11.)
- [Pre17] Thomas Prest. Sharper bounds in lattice-based cryptography using the Rényi divergence. In Tsuyoshi Takagi and Thomas Peyrin, editors, *Advances in Cryptology – ASIACRYPT 2017, Part I*, volume 10624 of *Lecture Notes in Computer Science*, pages 347–374, Hong Kong, China, December 3–7, 2017. Springer, Cham, Switzerland. doi:10.1007/978-3-319-70694-8_13. (Cited on pages 4, 6, 7, 12, 13, 19, 38, and 40.)
- [PS05] Thomas Pornin and Julien P. Stern. Digital signatures do not guarantee exclusive ownership. In John Ioannidis, Angelos Keromytis, and Moti Yung, editors, *ACNS 05: 3rd International Conference on Applied Cryptography and Network Security*, volume 3531 of *Lecture Notes in Computer Science*, pages 138–150, New York, NY, USA, June 7–10, 2005. Springer Berlin Heidelberg, Germany. doi:10.1007/11496137_10. (Cited on page 17.)
- [Rén61] Alfréd Rényi. On measures of entropy and information. *Proc. 4th Berkeley Symp. Math. Stat. Probab.* 1, 547–561 (1961)., 1961. (Cited on page 12.)
- [Rob23] Damien Robert. Breaking SIDH in polynomial time. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023, Part V*, volume 14008 of *Lecture Notes in Computer Science*, pages 472–503, Lyon, France, April 23–27, 2023. Springer, Cham, Switzerland. doi:10.1007/978-3-031-30589-4_17. (Cited on page 3.)
- [RSW18] Miruna Rosca, Damien Stehlé, and Alexandre Wallet. On the ring-LWE and polynomial-LWE problems. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018, Part I*, volume 10820 of *Lecture Notes in Computer Science*, pages 146–173, Tel Aviv, Israel, April 29 – May 3, 2018. Springer, Cham, Switzerland. doi:10.1007/978-3-319-78381-9_6. (Cited on page 3.)
- [SAB⁺22] Peter Schwabe, Roberto Avanzi, Joppe Bos, Léo Ducas, Eike Kiltz, Tancrede 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>. (Cited on page 3.)
- [SS11] Damien Stehlé and Ron Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In Kenneth G. Paterson, editor, *Advances in Cryptology – EUROCRYPT 2011*, volume 6632 of *Lecture Notes in Computer Science*, pages 27–47, Tallinn, Estonia, May 15–19, 2011. Springer Berlin Heidelberg, Germany. doi:10.1007/978-3-642-20465-4_4. (Cited on pages 3 and 6.)
- [TT15] Katsuyuki Takashima and Atsushi Takayasu. Tighter security for efficient lattice cryptography via the Rényi divergence of optimized orders. In Man Ho Au and Atsuko Miyaji, editors, *ProvSec 2015: 9th International Conference on Provable Security*, volume 9451 of *Lecture Notes in Computer Science*, pages 412–431, Kanazawa, Japan, November 24–26, 2015. Springer, Cham, Switzerland. doi:10.1007/978-3-319-26059-4_23. (Cited on pages 3, 15, and 19.)
- [vEH14] Tim van Erven and Peter Harremos. Rényi divergence and kullback-leibler divergence. *IEEE Transactions on Information Theory*, 60(7):3797–3820, 2014. doi:10.1109/TIT.2014.2320500. (Cited on page 26.)
- [YJW23] Yang Yu, Huiwen Jia, and Xiaoyun Wang. Compact lattice gadget and its applications to hash-and-sign signatures. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023, Part V*, volume 14085 of *Lecture Notes in Computer Science*, pages 390–420, Santa Barbara, CA, USA, August 20–24, 2023. Springer, Cham, Switzerland. doi:10.1007/978-3-031-38554-4_13. (Cited on pages 3 and 15.)

A Additional Preliminaries

Lemma 13 (Multiplicativity [LSS14, Lem. 4.1]). Let $a \in (1, \infty)$. Let $\mathcal{P}$ and $\mathcal{Q}$ denote distributions of a pair of random variables (Y_1, Y_2) . Also, for $i \in \{1, 2\}$ let $\mathcal{P}_i$ and $\mathcal{Q}_i$ be the marginal distribution of Y_i under $\mathcal{P}$ and $\mathcal{Q}$, respectively. Then if Y_1 and Y_2 are independent, $R_a(\mathcal{P} \parallel \mathcal{Q}) = R_a(\mathcal{P}_1 \parallel \mathcal{Q}_1) \cdot R_a(\mathcal{P}_2 \parallel \mathcal{Q}_2)$.

Lemma 14 (Probability Preservation [LSS14, Lem. 4.1]). Let $a \in (1, \infty)$ and $E \subseteq \text{sup}(\mathcal{Q})$ be an arbitrary event. Then,

$$\begin{aligned} \mathcal{Q}(E) &\geq \mathcal{P}(E)^{\frac{a}{a-1}} / R_a(\mathcal{P} \parallel \mathcal{Q}) \\ \mathcal{Q}(E) &\geq \mathcal{P}(E) / R_\infty(\mathcal{P} \parallel \mathcal{Q}). \end{aligned}$$

Lemma 15 (Data Processing Inequality [vEH14, Thm. 9]). Let $\alpha \in (1, \infty)$. For any function f , where $\mathcal{P}^f$ (respectively $\mathcal{Q}^f$) denotes the distribution of $f(x)$ induced by sampling $x \leftarrow \mathcal{P}$ (respectively $x \leftarrow \mathcal{Q}$), $R_a(\mathcal{P}^f \parallel \mathcal{Q}^f) \leq R_a(\mathcal{P} \parallel \mathcal{Q})$.

We use the following bound on the Rényi Divergence for Dependent Random Variables from [HPRR20].

Lemma 16 (Rényi Divergence for Dependent Random Variables [HPRR20, Prop. 4]). Let $\mathcal{P}$ and $\mathcal{Q}$ denote two distributions of an N -tuple of random variables $(X_i)_{i \leq N}$. For each $0 \leq i < N$, let $\mathcal{P}_i$ (resp. $\mathcal{Q}_i$) denote the marginal distribution of X_i , and let $\mathcal{P}_{i|<i}(\cdot \mid X_{<i})$ represent the conditional distribution of X_i given the values of the preceding variables $(X_0, \dots, X_{i-1}) = X_{<i}$. Let $a > 1$ and suppose that for every $0 \leq i < N$, there exists a constant $r_{a,i} \geq 1$ such that for every i -tuple $X_{<i}$ in the support of $\mathcal{Q}$ restricted to its first i variables,

$$\mathcal{R}_a(\mathcal{P}_{i|X_{<i}} \parallel \mathcal{Q}_{i|X_{<i}}) \leq r_{a,i}.$$

Then,

$$R_a(\mathcal{P} \parallel \mathcal{Q}) \leq \prod_{i < N} r_{a,i}.$$

B Proofs for Section 2 and Section 3

B.1 Proof of Lemma 8

Lemma 8 (Rényi Divergence of Gaussian Sample over Λ/Λ' (adapted from [GPV08, Cor. 2.8])). Let Λ, Λ' be n -dimensional full-rank lattices with $\Lambda' \subseteq \Lambda$. Then for any $a \in (1, \infty)$, $\epsilon \in (0, \frac{1}{2})$, any $s \geq \eta_\epsilon(\Lambda')$, and any $c \in \mathbb{R}^n$,

$$R_a(\mathcal{U}(\Lambda/\Lambda') \parallel \mathcal{D}_{\Lambda/\Lambda', s, c}) \lesssim 1 + \frac{2a\epsilon^2}{(1-\epsilon)^2}.$$

Proof. Much of the proof follows from [GPV08, Cor. 2.8], but for completeness and verifiability, we have fully proved these adaptations. The quotient group Λ/Λ' is defined as the additive group of cosets $x + \Lambda'$, $x \in \Lambda$. Sampling from a discrete Gaussian over this quotient group we obtain that for any $x \in \Lambda$

$$\mathcal{D}_{\Lambda/\Lambda', s, c}(x) = \frac{\rho_{s, c}(x + \Lambda')}{\rho_{s, c}(\Lambda)}.$$

By assumption $\Lambda' \subseteq \Lambda$ which implies $\eta_\epsilon(\Lambda) \leq \eta_\epsilon(\Lambda') \leq s$. Therefore, we can apply Lemma 5 and get

$$\rho_{s, c}(\Lambda) \in [1 - \epsilon, 1 + \epsilon] \cdot \frac{s^n}{\det(\Lambda)}.$$

Again, since $s \geq \eta_\epsilon(\Lambda')$

$$\rho_{s, c}(x + \Lambda') \in [1 - \epsilon, 1 + \epsilon] \cdot \frac{s^n}{\det(\Lambda')}.$$

Combining these results yields

$$\mathcal{D}_{\mathbf{\Lambda}/\mathbf{\Lambda}',s,c} \in \left[\frac{1-\epsilon}{1+\epsilon}, \frac{1+\epsilon}{1-\epsilon} \right] \cdot \frac{\det(\mathbf{\Lambda})}{\det(\mathbf{\Lambda}')}.$$

Since $\mathbf{\Lambda}$ and $\mathbf{\Lambda}'$ are full rank, their spans are the same ($\mathbb{R}^n$) and hence the size of their quotient group $\mathbf{\Lambda}/\mathbf{\Lambda}'$ is finite. Therefore, by [DD18, Lem. 10] we get that $|\mathbf{\Lambda}/\mathbf{\Lambda}'| = \frac{\det(\mathbf{\Lambda}')}{\det(\mathbf{\Lambda})}$. Computing the relative error between the Gaussian distribution and the uniform distribution $\mathcal{U}(\mathbf{\Lambda}/\mathbf{\Lambda}')(x) = \frac{1}{|\mathbf{\Lambda}/\mathbf{\Lambda}'|}$ gives

$$\frac{\mathcal{U}(|\mathbf{\Lambda}/\mathbf{\Lambda}'|)(x)}{\mathcal{D}_{\mathbf{\Lambda}/\mathbf{\Lambda}',s,c}(x)} \in \left[\frac{1-\epsilon}{1+\epsilon}, \frac{1+\epsilon}{1-\epsilon} \right] = \left[1 - \frac{2\epsilon}{1-\epsilon}, 1 + \frac{2\epsilon}{1-\epsilon} \right].$$

Applying Lemma 6 with $\delta = \frac{2\epsilon}{1-\epsilon}$, we obtain

$$R_a(\mathcal{U}(\mathbf{\Lambda}/\mathbf{\Lambda}') \parallel \mathcal{D}_{\mathbf{\Lambda}/\mathbf{\Lambda}',s,c}) \lesssim 1 + \frac{2a\epsilon^2}{(1-\epsilon)^2}.$$

This completes the proof. ■

B.2 Proof of Lemma 9

Lemma 9 (Rényi divergence (adapted from [GPV08, Lem 5.2])). If the columns of $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ generate $\mathbb{Z}_q^n$, $a \in (1, \infty)$, $\epsilon \in (0, \frac{1}{2})$, and $s \geq \eta_\epsilon(\mathbf{\Lambda}^\perp(\mathbf{A}))$; then for $e \sim \mathcal{D}_{\mathbb{Z}^m,s}$, the distribution $\mathcal{P} = \mathcal{U}(\mathbb{Z}_q^n)$, and the distribution $\mathcal{Q}$ of the syndromes $u = \mathbf{A}e \bmod q$, it holds that

$$R_a(\mathcal{P} \parallel \mathcal{Q}) \lesssim 1 + \frac{2a\epsilon^2}{(1-\epsilon)^2}.$$

Proof. For simplicity we denote $\mathbf{\Lambda}^\perp = \mathbf{\Lambda}^\perp(\mathbf{A})$. By assumption the set of all syndromes of $\mathbf{A}$ equals $\mathbb{Z}_q^n$, i.e. $\{\mathbf{A}e \bmod q \mid e \in \mathbb{Z}^m\} = \mathbb{Z}_q^n$. Consider the quotient group $(\mathbb{Z}^m/\mathbf{\Lambda}^\perp)$ which is defined as the group of all cosets, i.e. $\{e + \mathbf{\Lambda}^\perp \mid e \in \mathbb{Z}^m\}$. This quotient group is isomorphic to the set of syndromes of $\mathbf{A}$ via the mapping $e + \mathbf{\Lambda}^\perp \mapsto \mathbf{A}e \bmod q$, where $e \in \mathbb{Z}^m$. Hence, we have $\mathcal{P} \simeq \mathcal{U}(\mathbb{Z}^m/\mathbf{\Lambda}^\perp)$. Further, the distribution $\mathcal{D}_{\mathbb{Z}^m/\mathbf{\Lambda}^\perp,s} = \mathcal{D}_{\mathbb{Z}^m,s} \bmod \mathbf{\Lambda}^\perp$ is the distribution of $e \sim \mathcal{D}_{\mathbb{Z}^m,s}$ reduced modulo $\mathbf{\Lambda}^\perp$. That is, the coset $e + \mathbf{\Lambda}^\perp$ for $e \sim \mathcal{D}_{\mathbb{Z}^m,s}$. Applying the above isomorphism, this distribution is isomorphic to distribution $\mathcal{Q}$. Finally we can apply Lemma 8 with $\mathbf{\Lambda} = \mathbb{Z}^m$, $\mathbf{\Lambda}' = \mathbf{\Lambda}^\perp$ and $c = 0$ to obtain the claim. ■

B.3 Proof of Corollary 2

Corollary 2 (Rényi uniformity for NTRU). Let q be prime, $\mathbf{h} \in \mathcal{R}_q \setminus \{\mathbf{0}\}$, $a \in (1, \infty)$, $\epsilon \in (0, \frac{1}{2})$, $s \geq \eta_\epsilon(\mathbf{\Lambda}_{\mathbf{h},q})$, $\mathcal{P} = \mathcal{U}(\mathcal{R}_q)$, and $\mathcal{Q}$ the distribution of $\mathbf{u} + \mathbf{v} \cdot \mathbf{h} \bmod q$ where $\mathbf{u}, \mathbf{v} \sim \mathcal{D}_{\mathcal{R},s}$. Then it holds that

$$R_a(\mathcal{P} \parallel \mathcal{Q}) \lesssim 1 + \frac{2a\epsilon^2}{(1-\epsilon)^2}.$$

Proof. Elements in $\mathcal{R}$ are polynomials of degree n that can be described via their anticirculant matrix $\mathcal{A}(\cdot) \in \mathbb{Z}^{n \times n}$. For q prime and $\mathbf{h} \in \mathcal{R}_q \setminus \{\mathbf{0}\}$, we consider matrix $\mathbf{A} = [I_N | \mathcal{A}(\mathbf{h})] \in \mathbb{Z}^{n \times 2n}$ that defines the NTRU lattice $\mathbf{\Lambda}_{\mathbf{h},q} = \mathbf{\Lambda}^\perp(\mathbf{A})$. By Lemma 1 the anticirculant matrices with matrix addition and multiplication form a ring that is isomorphic to $\mathcal{R}$. In particular, this holds for the anticirculant of samples $e = (e_1, e_2)$ with $e_i \sim \mathcal{D}_{\mathbb{Z}^n,s}$ and $(\mathbf{u}, \mathbf{v})$ with $\mathbf{u}, \mathbf{v} \sim \mathcal{D}_{\mathcal{R},s}$ as well as for the resulting distributions $\mathbf{A} \cdot \mathcal{A}(e) \bmod q$ and the distribution of $\mathbf{z}$ such that $\mathcal{A}(\mathbf{z}) = \mathbf{A} \begin{bmatrix} \mathcal{A}(\mathbf{u}) \\ \mathcal{A}(\mathbf{v}) \end{bmatrix} = \mathcal{A}(\mathbf{u}) + \mathcal{A}(\mathbf{h}) \cdot \mathcal{A}(\mathbf{v}) \bmod q$. The latter distribution is equivalent to $\mathcal{Q}$. Finally, due to its special structure with identity I_N on the left, $\mathbf{A}$ generates $\mathbb{Z}_q^n$ such that we can apply Lemma 9 to conclude the proof. ■

B.4 Proof of Lemma 10

Lemma 10 (Relative Error for Tailbounds). Let $\mathcal{P}$ and $\mathcal{Q}$ be two distributions with $\text{sup}(\mathcal{P}) = \text{sup}(\mathcal{Q}) = \mathbb{Z}^n$ and $\delta_{RE}(\mathcal{P}, \mathcal{Q}) = \delta$. Then for any $\beta \geq 0$,

$$\Pr_{x \leftarrow \mathcal{P}}[\|x\|_2 > \beta] \leq \Pr_{x \leftarrow \mathcal{Q}}[\|x\|_2 > \beta] \cdot (1 + \delta).$$

Proof. We can use the relative error to upper bound the Rényi divergence of order ∞ :

$$R_\infty(\mathcal{P} \parallel \mathcal{Q}) = \max_{x \in \text{sup}(\mathcal{P})} \frac{\mathcal{P}(x)}{\mathcal{Q}(x)} \leq (1 + \delta).$$

Applying the probability preservation for R_∞ (Lemma 14) we obtain

$$\Pr_{x \leftarrow \mathcal{Q}}[\|x\|_2 > \beta] \geq \frac{\Pr_{x \leftarrow \mathcal{P}}[\|x\|_2 > \beta]}{R_\infty(\mathcal{P} \parallel \mathcal{Q})} \geq \Pr_{x \leftarrow \mathcal{P}}[\|x\|_2 > \beta] / (1 + \delta).$$

■

B.5 Proof of Lemma 11

Lemma 11 (Optimal Rényi Order). For $\lambda \in \mathbb{N}$, let $\mathcal{E}_1, \mathcal{E}_2$ be two events such that $\Pr[\mathcal{E}_1] \geq 2^{-\lambda}$. Assume that for any $Q \in \mathbb{N}$, $a \in (1, \infty)$, and $R_a \in [1, \infty)$ it holds that

$$\Pr[\mathcal{E}_2] \leq R_a^Q \cdot \Pr[\mathcal{E}_1]^{\frac{a-1}{a}}.$$

Then

$$-\log(\Pr[\mathcal{E}_2]) \geq -\log(\Pr[\mathcal{E}_1]) - \min_{a > 1} \left\{ Q \log R_a + \frac{\lambda}{a} \right\}.$$

Proof. By assumption it holds that $\Pr[\mathcal{E}_1] \geq 2^{-\lambda}$. Minimising for $a > 1$ yields

$$\begin{aligned} \Pr[\mathcal{E}_2] &\leq \min_{a > 1} \left\{ R_a^Q \cdot \Pr[\mathcal{E}_1]^{\frac{a-1}{a}} \right\} = \min_{a > 1} \left\{ R_a^Q \cdot \Pr[\mathcal{E}_1]^{-1/a} \right\} \cdot \Pr[\mathcal{E}_1] \\ &\leq \min_{a > 1} \left\{ R_a^Q \cdot 2^{\lambda/a} \right\} \cdot \Pr[\mathcal{E}_1]. \end{aligned}$$

In other words, this gives at least

$$-\log(\Pr[\mathcal{E}_1]) - \min_{a > 1} \left\{ Q \log R_a + \frac{\lambda}{a} \right\}$$

bits success probability for $\mathcal{E}_2$. ■

C Proof of Theorem 1

Theorem 1 (Unforgeability). For any adversary $\mathbf{A}$ against the **UF-CMA** security of COREFALCON^+ (Figure 6) running in time $t_{\mathbf{A}}$, making at most Q_s signing queries and $Q_{\mathbf{H}}$ random oracle queries, there exists an adversary $\mathbf{B}$ against $(Q_{\mathbf{H}} + 1)$ -**R-ISIS** running in time $t_{\mathbf{B}} \approx t_{\mathbf{A}}$ such that for all $C_s \in \mathbb{N}^{\geq 1}$ and $a_u, a_p \in \mathbb{R}^{>1}$ it holds

$$\begin{aligned} \text{Adv}_{\text{COREFALCON}^+, \mathbf{A}}^{Q_s\text{-}\mathbf{UF-CMA}} &\leq \left(r_u^{C_s} \cdot \left(r_p^{C_s} \cdot \text{Adv}_{q, \alpha, \beta, \mathbf{B}}^{(Q_{\mathbf{H}}+1)\text{-}\mathbf{R-ISIS}} \right)^{\frac{a_p-1}{a_p}} \right)^{\frac{a_u-1}{a_u}} \\ &\quad + \sum_{i=0}^{Q_s} \binom{C_s}{i} (1 - p_{\text{PreSmp}, \beta})^{C_s-i} (p_{\text{PreSmp}, \beta})^i + \frac{C_s(Q_{\mathbf{H}} + C_s)}{2^k}, \end{aligned}$$

where

$p_{\text{PreSmp},\beta} := \min_{(B,\cdot) \in \text{TpdGen}, c \in \mathcal{R}_q} \Pr_{(s_1, s_2) \leftarrow \text{PreSmp}(B, s, (c, 0))} [\|s_1, s_2\|_2 \leq \beta],$
 $r_u = \max_{(\cdot, h) \in \text{TpdGen}} R_{a_u}(\mathcal{P} \parallel \mathcal{Q}_h)$ with $\mathcal{P} = \mathcal{U}(\mathcal{R}_q)$ and $\mathcal{Q}_h$ the distribution of $s_1 + s_2 \cdot h \pmod q$, where $s_1, s_2 \sim \mathcal{D}_{\mathcal{R}, s}$,
 $r_p = \max_{(B,\cdot) \in \text{TpdGen}, c \in \mathcal{R}_q} R_{a_p}(\text{PreSmp}(B, s, (c, 0)) \parallel \mathcal{D}_{\Lambda, s})$ with $\Lambda = \Lambda(B)_{(c, 0)}$.

Consider the sequence of games depicted in Figure 7.

Game G_0 . This is the unforgeability game for COREFALCON^+ so by definition we have

$$\Pr[G_0^A \Rightarrow 1] = \text{Adv}_{\text{COREFALCON}^+, A}^{Q_s\text{-UF-CMA}}.$$

Games $G_0 - G_5$		Oracle $H(pk, r, m)$
01 $\mathcal{H}, \mathcal{Q} \leftarrow \emptyset$		19 if $\exists c : (c, pk, r, m) \in \mathcal{H}$
02 $\text{cnt} := 0$		20 return c
03 $(B, h) \leftarrow \text{Gen}$		21 $c \leftarrow \mathcal{R}_q$
04 $(m^*, \sigma^*) \leftarrow \text{A}^{\text{Sgn}(\cdot), H(\cdot, \cdot, \cdot)}(h)$		22 $\mathcal{H} \leftarrow \mathcal{H} \cup \{(c, pk, r, m)\}$
05 return $\llbracket \text{Ver}(h, m^*, \sigma^*) = 1 \wedge (m^*, \cdot) \notin \mathcal{Q} \rrbracket$		23 return c
Oracle $\text{Sgn}(m)$		Oracle $H'(pk, r, m)$
06 repeat		24 $pk \rightarrow h$
07 $\text{cnt} \leftarrow \text{cnt} + 1$	// $G_1 - G_5$	25 if $\exists c : (c, h, r, m) \in \mathcal{H}$
08 if $\text{cnt} > C_s$	// $G_1 - G_5$	26 abort
09 abort	// $G_1 - G_5$	27 $c \leftarrow \mathcal{R}_q$
10 $r \leftarrow \{0, 1\}^k$		28 $(s_1, s_2) := (\perp, \perp)$
11 $c := H(h, r, m)$	// $G_0 - G_1$	29 $s_1, s_2 \leftarrow \mathcal{D}_{\mathcal{R}, s}$ // $G_3 - G_5$
12 $(c, s_1, s_2) := H'(h, r, m)$	// $G_2 - G_5$	30 $c := s_1 + s_2 \cdot h \pmod q$ // $G_3 - G_5$
13 $(s_1, s_2) \leftarrow \text{PreSmp}(B, s, (c, 0))$	// $G_0 - G_3$	31 $\mathcal{H} \leftarrow \mathcal{H} \cup \{(c, h, r, m)\}$
14 $(s_1, s_2) \leftarrow \mathcal{D}_{\Lambda(B)_{(c, 0), s}}$	// G_4	32 return (c, s_1, s_2)
15 until $\ s_1, s_2\ _2 \leq \beta$		
16 $\sigma := (r, s_2)$		
17 $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(m, \sigma)\}$		
18 return σ		

Figure 7. Games for the proof of Theorem 1.

Game G_1 . This game is identical to the previous one, except that it aborts if the overall number of sampled preimages in the signing oracle, i.e. including potential repetitions, exceeds threshold C_s .

Claim 1: For $p_{\text{PreSmp},\beta} := \min_{(B,\cdot) \in \text{TpdGen}, c \in \mathcal{R}_q} \Pr_{(s_1, s_2) \leftarrow \text{PreSmp}(B, s, (c, 0))} [\|s_1, s_2\|_2 \leq \beta]$ it holds that

$$|\Pr[G_0^A \Rightarrow 1] - \Pr[G_1^A \Rightarrow 1]| \leq \sum_{i=0}^{Q_s} \binom{C_s}{i} (1 - p_{\text{PreSmp},\beta})^{C_s-i} (p_{\text{PreSmp},\beta})^i.$$

Proof. To proof the claim, we model the experiment using a binomial distributed random variable $X \sim B(C_s, p_{\text{PreSmp},\beta})$, i.e. we have C_s Bernoulli trials and success probability $p_{\text{PreSmp},\beta}$. A trial corresponds to sampling a preimage using PreSmp in the signing oracle and the trial succeeds if the norm is sufficiently small, i.e. $\|s_1, s_2\|_2 \leq \beta$. Hence, the random variable, counting the overall number of successes in the Bernoulli trials, tells us the number of signing queries we are able to answer. Since we need to answer Q_s signing queries, we are interested in the CDF for value Q_s , i.e. $\Pr[X \leq Q_s]$ which is exactly the claim. ■

Game $\mathbf{G}_2$. This game is identical to the previous one, except that it aborts during a signing oracle query if there already exists a query to the random oracle for the same public key, salt r , and message m as the output signature. To ease the depiction in further hybrids, we define a new RO $\mathbf{H}'$ maintaining the same set $\mathcal{H}$ as $\mathbf{H}$ but aborting in case of a query on the same input as a previous query. Oracle $\mathbf{H}'$ is then called within the signing oracle instead of $\mathbf{H}$.

Claim 2: $|\Pr[\mathbf{G}_1^A \Rightarrow 1] - \Pr[\mathbf{G}_2^A \Rightarrow 1]| \leq \frac{C_s(C_s + Q_H)}{2^k}$.

Proof of Claim. The salt r is chosen uniformly at random from $\{0, 1\}^k$ for each RO query during a signing query. The total number of elements in $\mathcal{H}$ is at most $C_s + Q_H$, as at most one element is added per query to $\mathbf{H}$ (or $\mathbf{H}'$). Thus, the probability that the freshly chosen salt was part of a previous query is at most $\frac{C_s + Q_H}{2^k}$. For C_s queries to the internal oracle $\mathbf{H}'$, we obtain the claimed bound. $\square$

Game $\mathbf{G}_3$. This game is the same as the previous one, except that random oracle $\mathbf{H}'$ no longer returns a uniformly random element $\mathbf{c} \leftarrow^{\$} \mathcal{R}_q$. Instead, it computes $\mathbf{c}$ as follows: It samples two elements $\mathbf{s}_1, \mathbf{s}_2$ from a Gaussian distribution $\mathcal{D}_{\mathcal{R}, s}$ with standard deviation s over ring $\mathcal{R}$. Then, $\mathbf{c}$ is computed as $\mathbf{c} = \mathbf{s}_1 + \mathbf{s}_2 \cdot \mathbf{h}$ mod q , where $\mathbf{h}$ is the public key. For future use, $\mathbf{s}_1, \mathbf{s}_2$ are returned together with the RO output (note that $\mathbf{H}'$ cannot be called directly by the adversary).

Claim 3: For $\mathbf{h} \in \mathcal{R}$, let $\mathcal{P} := \mathcal{U}(\mathcal{R}_q)$ and $\mathcal{Q}_{\mathbf{h}}$ be the distribution of $\mathbf{s}_1 + \mathbf{s}_2 \cdot \mathbf{h}$ mod q where $\mathbf{s}_1, \mathbf{s}_2 \leftarrow^{\$} \mathcal{D}_{\mathcal{R}, s}$. Then, for any $a_u \in (1, \infty)$,

$$\Pr[\mathbf{G}_2^A \Rightarrow 1] \leq \left(\max_{(\cdot, \mathbf{h}) \in \text{TpdGen}} R_{a_u}(\mathcal{P} \parallel \mathcal{Q}_{\mathbf{h}})^{C_s} \cdot \Pr[\mathbf{G}_3^A \Rightarrow 1] \right)^{\frac{a_u - 1}{a_u}}.$$

Proof of Claim. We define two underlying distributions for a $(Q + 1)$ -tuple of random variables $(\mathbf{c}_0 = (\mathbf{B}, \mathbf{h}), \mathbf{c}_1, \dots, \mathbf{c}_Q)$.

$\bar{\mathcal{P}}$ $(\mathbf{B}, \mathbf{h}) \leftarrow^{\$} \text{TpdGen}$ for $i \in [Q]$ $\mathbf{c}_i \leftarrow^{\$} \mathcal{R}_q$ return $((\mathbf{B}, \mathbf{h}), \mathbf{c}_1, \dots, \mathbf{c}_Q)$	$\bar{\mathcal{Q}}$ $(\mathbf{B}, \mathbf{h}) \leftarrow^{\$} \text{TpdGen}$ for $i \in [Q]$ $(\mathbf{s}_1, \mathbf{s}_2) \leftarrow \mathcal{D}_{\mathcal{R}, s}$ $\mathbf{c}_i := \mathbf{s}_1 + \mathbf{s}_2 \cdot \mathbf{h} \text{ mod } q$ return $((\mathbf{B}, \mathbf{h}), \mathbf{c}_1, \dots, \mathbf{c}_Q)$
--	---

These distributions describe the underlying distributions of $\mathbf{G}_2$ and $\mathbf{G}_3$. By the data processing inequality (Lemma 15) it holds that, for any $a \in (1, \infty)$,

$$R_a(\mathbf{G}_2 \parallel \mathbf{G}_3) \leq R_a(\bar{\mathcal{P}} \parallel \bar{\mathcal{Q}}). \quad (1)$$

Let the marginal distribution of $\mathbf{c}_i$ be denoted by $\bar{\mathcal{P}}_i$ ($\bar{\mathcal{Q}}_i$ resp.) and the distribution of $\mathbf{c}_i$ conditioned on $\mathbf{c}_{<i} = (\mathbf{c}_0, \dots, \mathbf{c}_{i-1})$ as $\bar{\mathcal{P}}_{i|\mathbf{c}_{<i}}$ ($\bar{\mathcal{Q}}_{i|\mathbf{c}_{<i}}$ resp.). Since the distribution of $\mathbf{c}_0 = (\mathbf{B}, \mathbf{h})$ is the same for $\bar{\mathcal{P}}$ and $\bar{\mathcal{Q}}$, it holds that

$$R_a(\bar{\mathcal{P}}_0 \parallel \bar{\mathcal{Q}}_0) = 1.$$

For the conditional distributions, note that random variable $\mathbf{c}_i$ is independent of the previous random variables $\mathbf{c}_1, \dots, \mathbf{c}_{i-1}$. However, $\mathbf{c}_i$ might depend on $\mathbf{h}$ and thus on random variable $\mathbf{c}_0 = (\mathbf{B}, \mathbf{h})$. Hence for all $i \in [Q + 1]$,

$$\begin{aligned} R_a(\bar{\mathcal{P}}_{i|\mathbf{c}_{<i}} \parallel \bar{\mathcal{Q}}_{i|\mathbf{c}_{<i}}) &= R_a(\bar{\mathcal{P}}_{i|(\mathbf{c}_0, \dots, \mathbf{c}_{i-1})} \parallel \bar{\mathcal{Q}}_{i|(\mathbf{c}_0, \dots, \mathbf{c}_{i-1})}) \\ &\leq \max_{(\mathbf{B}, \mathbf{h}) \in \text{TpdGen}} R_a(\bar{\mathcal{P}}_{i|((\mathbf{B}, \mathbf{h}), \mathbf{c}_1, \dots, \mathbf{c}_{i-1})} \parallel \bar{\mathcal{Q}}_{i|((\mathbf{B}, \mathbf{h}), \mathbf{c}_1, \dots, \mathbf{c}_{i-1})}) \\ &= \max_{(\cdot, \mathbf{h}) \in \text{TpdGen}} R_a(\mathcal{P} \parallel \mathcal{Q}_{\mathbf{h}}), \end{aligned}$$

where $\mathcal{P} := \mathcal{U}(\mathcal{R}_q)$ and $\mathcal{Q}_h$ the distribution of $\mathbf{s}_1 + \mathbf{s}_2 \cdot \mathbf{h} \bmod q$ where $\mathbf{s}_1, \mathbf{s}_2 \xleftarrow{\$} \mathcal{D}_{\mathcal{R},s}$. Note that $\mathbf{h}$ does not occur in distribution $\mathcal{P}$ because the individual random variables $\mathbf{c}_i$ (for $i \geq 1$) are independent of $\mathbf{h}$.

By Lemma 16 it follows

$$R_a(\bar{\mathcal{P}} \parallel \bar{\mathcal{Q}}) \leq \max_{(\cdot, \mathbf{h}) \in \text{TpGen}} R_a(\mathcal{P} \parallel \mathcal{Q}_h)^Q. \quad (2)$$

Combining probability preservation (Lemma 14) with Equation (1) and Equation (2), we obtain

$$\Pr[\mathbf{G}_3^A \Rightarrow 1] \geq \frac{\Pr[\mathbf{G}_2^A \Rightarrow 1]^{\frac{a}{a-1}}}{R_a(\mathbf{G}_2 \parallel \mathbf{G}_3)} \geq \frac{\Pr[\mathbf{G}_2^A \Rightarrow 1]^{\frac{a}{a-1}}}{\max_{(\cdot, \mathbf{h}) \in \text{TpGen}} R_a(\mathcal{P} \parallel \mathcal{Q}_h)^Q}.$$

The claim follows by setting $Q := C_s$ due to at most C_s queries from Sgn to $\mathbf{H}'$ in Line 12. $\square$

Game $\mathbf{G}_4$. This game is identical to the previous one except that the output of the preimage sampler $\text{PreSmp}(\mathbf{B}, s, (\mathbf{c}, \mathbf{0}))$ is replaced by a Gaussian over the lattice $\mathbf{\Lambda} = \mathbf{\Lambda}(\mathbf{B})_{(\mathbf{c}, \mathbf{0})}$, namely $\mathcal{D}_{\mathbf{\Lambda}, s}$.

Claim 4: For distributions $\text{PreSmp} := \text{PreSmp}(\mathbf{B}, s, (\mathbf{c}, \mathbf{0}))$, $\mathcal{D} := \mathcal{D}_{\mathbf{\Lambda}, s}$, and $a_p \in (1, \infty)$ it holds that

$$\Pr[\mathbf{G}_3^A \Rightarrow 1] \leq \max_{(\mathbf{B}, \cdot) \in \text{TpGen}, \mathbf{c} \in \mathcal{R}_q} \left(R_{a_p}(\text{PreSmp} \parallel \mathcal{D})^{C_s} \cdot \Pr[\mathbf{G}_4^A \Rightarrow 1] \right)^{\frac{a_p-1}{a_p}}.$$

Proof of Claim. The claim follows analogously to Game $\mathbf{G}_3$. $\square$

Game $\mathbf{G}_5$. This game is identical to the previous one except that preimages $\mathbf{s}_1, \mathbf{s}_2$ are not sampled from a Gaussian distribution over the lattice shifted by $(\mathbf{c}, \mathbf{0})$ as before. Instead, the preimages of $\mathbf{c}$ that were sampled in $\mathbf{H}'$ are used.

Claim 5: $\Pr[\mathbf{G}_4^A \Rightarrow 1] = \Pr[\mathbf{G}_5^A \Rightarrow 1]$.

Proof of Claim. We need to show that the distributions of the games are the same. The RO output $\mathbf{c}$ is the same in both games. In $\mathbf{G}_4$, the signing oracle outputs $(\mathbf{s}_1, \mathbf{s}_2) \sim \mathcal{D}_{\mathbf{\Lambda}(\mathbf{B})_{(\mathbf{c}, \mathbf{0})}, s}$. Since $\mathbf{\Lambda}(\mathbf{B})$ is the NTRU lattice shifted by $(\mathbf{c}, \mathbf{0})$, the output is distributed according to a Gaussian $\mathcal{D}_{\mathcal{R}, s}$ conditioned on $\mathbf{s}_1 + \mathbf{s}_2 \cdot \mathbf{h} = \mathbf{c} \bmod q$. The output distribution in Game $\mathbf{G}_5$ is a Gaussian $\mathcal{D}_{\mathcal{R}, s}$ as well where the condition $\mathbf{s}_1 + \mathbf{s}_2 \cdot \mathbf{h} = \mathbf{c} \bmod q$ is fulfilled by construction (Line 30). $\square$

Reduction from $\mathcal{R}$ -ISIS. Claim 6: There exists an adversary $\mathbf{B}$ against $(Q_H + 1)$ - $\mathcal{R}$ -ISIS such that

$$\Pr[\mathbf{G}_5^A \Rightarrow 1] \leq \text{Adv}_{q, \alpha, \beta, \mathbf{B}}^{(Q_H+1)\text{-}\mathcal{R}\text{-ISIS}}.$$

Proof of Claim. Adversary $\mathbf{B}$ is formally constructed in Figure 8. Due to the changes in the previous games, adversary $\mathbf{B}$ can perfectly simulate the game for adversary $\mathbf{A}$ against $\mathbf{G}_5$ without having the secret key for $\mathbf{h}$. Further, $\mathbf{B}$ embeds their own targets in the queries to $\mathbf{H}$. Let us assume, that $\mathbf{A}$ wins $\mathbf{G}_5$, i.e. the forgery verifies and $(m^*, \cdot)$ was not queried to Sgn before. This implies that there exists an i^* such that $\hat{\mathbf{c}}_{i^*} = \mathbf{c}^*$ because if $\mathbf{A}$ wins the game, the challenge RO output $\mathbf{c}^*$ is equal to one of $\mathbf{B}$'s targets (that is exactly $\hat{\mathbf{c}}_{i^*}$) or to a signing query. If it corresponds to a signing query, there is no way that adversary $\mathbf{A}$ can win the game due to the freshness condition $(m^*, \cdot) \notin \mathcal{Q}$. Hence, Line 06 ensures the first winning condition of $\mathbf{B}$, which is $\mathbf{s}_1^* + \mathbf{s}_2^* \cdot \mathbf{h} = \hat{\mathbf{c}} \bmod q$. Further, the norm bound from $\mathbf{A}$ directly translates to the second winning condition, i.e. $\|(\mathbf{s}_1^*, \mathbf{s}_2^*)\|_2 \leq \beta$. $\square$

<u>B($h, \hat{c}_1, \dots, \hat{c}_{Q_H+1}$)</u>	<u>Oracle H(pk, r, m)</u>
01 $\mathcal{H}, \mathcal{Q} \leftarrow \emptyset$	10 if $\exists c : (c, pk, r, m) \in \mathcal{H}$
02 $\text{cnt}, \ell := 0$	11 return c
03 $(m^*, \sigma^*) \leftarrow \mathbb{A}^{\text{Sgn}(\cdot), \text{H}(\cdot, \cdot, \cdot)}(h)$	12 $\ell := \ell + 1$
04 parse $\sigma^* \rightarrow (r^*, s_2^*)$	13 $c := \hat{c}_\ell$ // embed challenge target
05 $c^* := \text{H}(h, r^*, m^*)$	14 $\mathcal{H} \leftarrow \mathcal{H} \cup \{(c, pk, r, m)\}$
06 $s_1^* := c^* - s_2^* \cdot h \pmod{q}$	15 return c
07 find $i^* : c^* = \hat{c}_{i^*}$	<u>Oracle H'(h, r, m)</u>
08 return (i^*, s_1^*, s_2^*)	16 return $\mathbf{G}_5.H'(h, r, m)$
<u>Oracle Sgn(m)</u>	
09 return $\mathbf{G}_5.\text{Sgn}(m)$	

Figure 8. Adversary B against t - $\mathcal{R}$ -ISIS for the proof of Theorem 1.

D Proof of Theorem 2

Theorem 2 (Strong Unforgeability). For any adversary A against the **SUF-CMA** security of COREFALCON^+ (Figure 6) running in time t_A , making at most Q_s signing queries and Q_H random oracle queries, there exist an adversary B against Q_s - $\mathcal{R}$ -SPISIS running in time $t_B \approx t_A$ such that for all $C_s \in \mathbb{N}^{\geq 1}$ and $a_p \in \mathbb{R}^{>1}$ it holds

$$\begin{aligned} \text{Adv}_{\text{COREFALCON}^+, A}^{Q_s\text{-SUF-CMA}} &\leq \text{Adv}_{\text{COREFALCON}^+, A}^{Q_s\text{-UF-CMA}} + \left(r_p^{C_s} \cdot \left(\text{Adv}_{q, \alpha, s, \beta, B}^{Q_s\text{-}\mathcal{R}\text{-SPISIS}} + p_{\text{binom}} \right) \right)^{\frac{a_p - 1}{a_p}} \\ &\quad + p_{\text{binom}} + \left(\frac{Q_s + 1}{2p_{\text{PreSmp}, \beta}^2} + \frac{2Q_H}{p_{\text{PreSmp}, \beta}} \right) Q_s 2^{-k}, \end{aligned}$$

where

$$\begin{aligned} p_{\text{PreSmp}, \beta} &:= \min_{(B, \cdot) \in \text{TpdGen}, c \in \mathcal{R}_q} \Pr_{(s_1, s_2) \leftarrow \mathbb{A}^{\text{PreSmp}}(B, s, (c, 0))} [\| (s_1, s_2) \|_2 \leq \beta], \\ p_{\text{binom}} &:= \sum_{i=0}^{Q_s} \binom{C_s}{i} (1 - p_{\text{PreSmp}, \beta})^{C_s - i} (p_{\text{PreSmp}, \beta})^i \\ r_p &= \max_{(B, \cdot) \in \text{TpdGen}, c \in \mathcal{R}_q} R_{a_p}(\text{PreSmp}(B, s, (c, 0)) \parallel \mathcal{D}_{\Lambda, s}) \text{ with } \Lambda = \Lambda(B)_{(c, 0)}. \end{aligned}$$

Proof. We prove the theorem by a sequence of games.

Game $\mathbf{G}_0$. We start with the **SUF-CMA** game for COREFALCON^+ :

$$\Pr[\mathbf{G}_0^A \Rightarrow 1] = \text{Adv}_{\text{COREFALCON}^+, A}^{Q_s\text{-SUF-CMA}}.$$

To prove the theorem, we distinguish between two kind of adversaries. One is a plain unforgeability adversary who returns a forgery which is the preimage of an RO output that corresponds to a previous query to the signing oracle. Such an adversary directly reduces to the bound of Theorem 1. The second adversary returns a forgery corresponding to a RO oracle query that was issued in the signing oracle.

Game $\mathbf{G}_1$. This is the same game as the previous one except that the adversary only outputs forgeries such that the associated RO was programmed during a signing query.

$$\left| \Pr[\mathbf{G}_0^A \Rightarrow 1] - \Pr[\mathbf{G}_1^A \Rightarrow 1] \right| \leq \text{Adv}_{\text{COREFALCON}^+, A}^{Q_s\text{-UF-CMA}}.$$

Games $G_0 - G_5$	Oracle $\text{Sgn}(m)$
01 $\mathcal{H}, \mathcal{Q} \leftarrow \emptyset$	13 repeat
02 $\text{cnt} := 0$	14 $\text{cnt} \leftarrow \text{cnt} + 1$ // $G_3 - G_4$
03 $(m^*, \sigma^*) \xleftarrow{\$} A^{\text{Sgn}(\cdot), H(\cdot, \cdot, \cdot)}(h)$	15 if $\text{cnt} > C_s$ // $G_3 - G_4$
04 parse $\sigma^* \rightarrow (r^*, s_2^*)$	16 abort // $G_3 - G_4$
05 $c^* := H(h, r^*, m^*)$	17 $r \xleftarrow{\$} \{0, 1\}^k$
06 $s_1^* := c^* - s_2^* \cdot h \pmod q$	18 $c := H(h, r, m)$ // $G_0 - G_1$
07 return $\llbracket \text{Ver}(h, m^*, \sigma^*) = 1 \wedge (m^*, \cdot) \notin \mathcal{Q} \rrbracket$	19 $c := H'(h, r, m)$ // $G_2 - G_5$
Oracle $H'(h, r, m)$ // $G_2 - G_5$	20 $(s_1, s_2) \xleftarrow{\$} \text{PreSmp}(B, s, (c, 0))$
08 flag := false	21 $(s_1, s_2) \xleftarrow{\$} \mathcal{D}_{\Lambda(B)(c, 0), s}$ // $G_4 - G_5$
09 if $\exists c : (c, h, r, m) \in \mathcal{H}$	22 until $\ (s_1, s_2)\ _2 \leq \beta$
10 $\text{flag} := \text{true}$	23 if flag = true // $G_2 - G_5$
11 $c \xleftarrow{\$} \mathcal{R}_q$	24 abort // $G_2 - G_5$
12 return c	25 $\mathcal{H} := \mathcal{H} \cup \{(c, h, r, m)\}$ // $G_2 - G_5$
	26 $\sigma := (r, s_2)$
	27 $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(m, \sigma)\}$
	28 return σ
	Oracle $H(pk, r, m)$
	29 if $\exists c : (c, pk, r, m) \in \mathcal{H}$
	30 return c
	31 $c \xleftarrow{\$} \mathcal{R}_q$
	32 $\mathcal{H} \leftarrow \mathcal{H} \cup \{(c, pk, r, m)\}$
	33 return c

Figure 9. Games $G_0 - G_5$ for the proof of Theorem 2.

Game G_2 . This is the same game as the previous one except that it aborts in the signing oracle if the random oracle was already queried on the same input before. We only consider the query corresponding to the signature that is eventually output by the signing oracle. Additionally, we only program the random oracle on that query and ignore the other repetitions before.

Claim 7: It holds that

$$\left| \Pr[G_1^A \Rightarrow 1] - \Pr[G_2^A \Rightarrow 1] \right| \leq \left(\frac{Q_s + 1}{2p_{\text{PreSmp}, \beta}^2} + \frac{2Q_H}{p_{\text{PreSmp}, \beta}} \right) Q_s 2^{-k}.$$

Proof of Claim. The claim is implicit in [BBD⁺23, Thm. 3], where it is separated into two steps. The first step is denoted by the difference between “Sign” and “Prog”, and the second one by the difference between “Prog” and “Trans”. $\square$

Game G_3 . This game is identical to the previous one, except that it aborts if the overall number of sampled preimages in the signing oracle, i.e. including potential repetitions, exceeds threshold C_s .

Claim 8: For $p_{\text{PreSmp}, \beta} := \min_{(B, \cdot) \in \text{TpdGen}, c \in \mathcal{R}_q} \Pr_{(s_1, s_2) \xleftarrow{\$} \text{PreSmp}(B, s, (c, 0))} [\|(s_1, s_2)\|_2 \leq \beta]$ it holds that

$$\left| \Pr[G_2^A \Rightarrow 1] - \Pr[G_3^A \Rightarrow 1] \right| \leq \sum_{i=0}^{Q_s} \binom{C_s}{i} (1 - p_{\text{PreSmp}, \beta})^{C_s - i} (p_{\text{PreSmp}, \beta})^i.$$

Proof. To proof the claim, we model the experiment using a binomial distributed random variable $X \sim B(C_s, p_{\text{PreSmp}, \beta})$, i.e. we have C_s Bernoulli trials and success probability $p_{\text{PreSmp}, \beta}$. A trial corresponds to

sampling a preimage using **PreSmp** in the signing oracle and the trial succeeds if the norm is sufficiently small, i.e. $\|(\mathbf{s}_1, \mathbf{s}_2)\|_2 \leq \beta$. Hence, the random variable, counting the overall number of successes in the Bernoulli trials, tells us the number of signing queries we are able to answer. Since we need to answer Q_s signing queries, we are interested in the CDF for value Q_s , i.e. $\Pr[X \leq Q_s]$ which is exactly the claim. ■

Game $\mathbf{G}_4$. This game is identical to the previous one except that the output of the preimage sampler $\text{PreSmp}(\mathbf{B}, s, (\mathbf{c}, \mathbf{0}))$ is replaced by a Gaussian over the lattice $\mathbf{\Lambda} = \mathbf{\Lambda}(\mathbf{B})_{(\mathbf{c}, \mathbf{0})}$, namely $\mathcal{D}_{\mathbf{\Lambda}, s}$.

Claim 9: For distributions $\text{PreSmp} := \text{PreSmp}(\mathbf{B}, s, (\mathbf{c}, \mathbf{0}))$, $\mathcal{D} := \mathcal{D}_{\mathbf{\Lambda}, s}$, and $a_p \in (1, \infty)$ it holds that

$$\Pr[\mathbf{G}_3^{\mathbf{A}} \Rightarrow 1] \leq \max_{(\mathbf{B}, \cdot) \in \text{TpdGen}, \mathbf{c} \in \mathcal{R}_q} \left(R_{a_p}(\text{PreSmp} \parallel \mathcal{D})^{C_s} \cdot \Pr[\mathbf{G}_4^{\mathbf{A}} \Rightarrow 1] \right)^{\frac{a_p - 1}{a_p}}.$$

Proof of Claim. The claim follows analogously to the proof of Theorem 1. □

Game $\mathbf{G}_5$. This game reverts the changes made in $\mathbf{G}_3$. Claim 10: For

$$p_{\text{PreSmp}, \beta} := \min_{(\mathbf{B}, \cdot) \in \text{TpdGen}, \mathbf{c} \in \mathcal{R}_q} \Pr_{(\mathbf{s}_1, \mathbf{s}_2) \leftarrow \text{PreSmp}(\mathbf{B}, s, (\mathbf{c}, \mathbf{0}))} [\|(\mathbf{s}_1, \mathbf{s}_2)\|_2 \leq \beta]$$

it holds that

$$\left| \Pr[\mathbf{G}_4^{\mathbf{A}} \Rightarrow 1] - \Pr[\mathbf{G}_5^{\mathbf{A}} \Rightarrow 1] \right| \leq \sum_{i=0}^{Q_s} \binom{C_s}{i} (1 - p_{\text{PreSmp}, \beta})^{C_s - i} (p_{\text{PreSmp}, \beta})^i.$$

Proof. To proof is analogous to the proof for $\mathbf{G}_3$. ■

Final reduction. We can reduce $\mathbf{G}_5$ to t - $\mathcal{R}$ -**SPISIS**.

Claim 11: There exists an adversary $\mathbf{C}$ against t - $\mathcal{R}$ -**SPISIS** such that

$$\Pr[\mathbf{G}_5^{\mathbf{A}} \Rightarrow 1] \leq \text{Adv}_{q, \alpha, s, \beta, \mathbf{B}}^{Q_s - \mathcal{R}\text{-SPISIS}}.$$

Proof of Claim. Reduction $\mathbf{B}$ is formally constructed in Figure 10.

The simulations of the random oracle and the signing oracle **Sgn** are perfect since the distribution of $\mathbf{B}$'s inputs exactly follow the distributions required in $\mathbf{G}_5$. This is the case because exactly one random oracle position is programmed during a signing query. $\mathbf{B}$ can answer the signing query because their preimages have a sufficiently small norm. If $\mathbf{A}$ wins, the reduction finds an index in Line 07 because we are only considering forgeries that correspond to a previous signing query and due to $\mathbf{A}$'s freshness condition, namely $(m^*, \sigma^*) \notin \mathcal{Q}$. Note that if $\mathbf{A}$ wins their game, all winning conditions of $\mathbf{C}$ are fulfilled. First, $(\mathbf{s}_1^*, \mathbf{s}_2^*)$ is a preimage of $\mathbf{c}_{\text{cnt}^*}$ (Line 06) which must have a norm of at most β if $\mathbf{A}$ wins their unforgeability game. Lastly, due to the check in Line 07, the output solution must be fresh. □

This completes the proof. ■

E Appendix for Section 5

<u>Adversary $B(h, \{(c_j, u_j, v_j)\}_{j \in [Q_s]})$</u>	<u>Oracle $\text{Sgn}(m)$</u>
01 $\mathcal{H}, \mathcal{Q} \leftarrow \emptyset$	09 $\text{cnt} \leftarrow \text{cnt} + 1$
02 $\text{cnt} := 0$	10 $r \xleftarrow{\$} \{0, 1\}^k$
03 $(m^*, \sigma^*) \xleftarrow{\$} A^{\text{Sgn}(\cdot), H(\cdot, \cdot, \cdot)}(h)$	11 $c := c_{\text{cnt}}$ // embed target
04 parse $\sigma^* \rightarrow (r^*, s_2^*)$	12 $(s_1, s_2) := (u_{\text{cnt}}, v_{\text{cnt}})$ // embed preimage
05 $c^* := H(h, r^*, m^*)$	13 $\mathcal{H} := \mathcal{H} \cup \{(c, h, r, m)\}$ // program RO
06 $s_1^* := c^* - s_2^* \cdot h \pmod{q}$	14 $\sigma := (r, s_2)$
07 find $\text{cnt}^* : c^* = c_{\text{cnt}^*} \wedge (s_1^*, s_2^*) \neq (u_{\text{cnt}^*}, v_{\text{cnt}^*})$	15 $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(m, \sigma)\}$
08 return $(\text{cnt}^*, s_1^*, s_2^*)$	16 return σ
	<u>Oracle $H(pk, r, m)$</u>
	17 if $\exists c : (c, pk, r, m) \in \mathcal{H}$
	18 return c
	19 $c \xleftarrow{\$} \mathcal{R}_q$
	20 $\mathcal{H} \leftarrow \mathcal{H} \cup \{(c, pk, r, m)\}$
	21 return c

Figure 10. Adversary C against t - $\mathcal{R}$ -SPISIS for the proof of Theorem 2 simulating G_5 .

```
sage: SIS.estimate.rough(SIS.Parameters(n=512,q=12289,length_bound=5833.93,norm=2,m=2*512))
lattice :: rop:  $\approx 2^{121.2}$ , red:  $\approx 2^{121.2}$ ,  $\delta$ : 1.003882,  $\beta$ : 415, d: 1024, tag: euclidean
sage: SIS.estimate.rough(SIS.Parameters(n=1024,q=12289,length_bound=8382.44,norm=2,m=2*1024))
lattice :: rop:  $\approx 2^{279.2}$ , red:  $\approx 2^{279.2}$ ,  $\delta$ : 1.002114,  $\beta$ : 956, d: 2048, tag: euclidean
```

Figure 11. SIS hardness estimates for ring dimension $n = 512$, $n = 1024$ and length bound β .

E.1 Proof of Lemma 12

Lemma 12 (Optimal C_s). For FALCON^+ -512 with $\lambda = 128$ it holds that,

$$\arg \min_{C_s} \left\{ C_s \left| \sum_{i=0}^{Q_s} \binom{C_s}{i} (1 - p_{\text{PreSmp}, \beta})^{C_s - i} (p_{\text{PreSmp}, \beta})^i \leq 2^\lambda \right. \right\} \lesssim 2^{64} + 2^{50},$$

and for FALCON^+ -1024 with $\lambda = 256$ it holds that

$$\arg \min_{C_s} \left\{ C_s \left| \sum_{i=0}^{Q_s} \binom{C_s}{i} (1 - p_{\text{PreSmp}, \beta})^{C_s - i} (p_{\text{PreSmp}, \beta})^i \leq 2^\lambda \right. \right\} \lesssim 2^{64} + 2^{36}.$$

Proof. First, we compute $p_{\text{PreSmp}, \beta}$ as follows.

$$\begin{aligned}
p_{\text{PreSmp},\beta} &:= \min_{\substack{c \in \mathcal{R}_q \\ (B, \cdot) \in \text{TpGen}}} \Pr_{(s_1, s_2) \leftarrow \mathbb{S}\text{PreSmp}(B, s, (c, 0))} [\| (s_1, s_2) \|_2 \leq \beta] \\
&= \min_{\substack{c \in \mathcal{R}_q \\ (B, \cdot) \in \text{sup}(\text{Gen})}} 1 - \Pr_{(s_1, s_2) \leftarrow \mathbb{S}\text{PreSmp}(B, s, (c, 0))} [\| (s_1, s_2) \|_2 > \beta] \\
&\gtrsim \min_{\substack{c \in \mathcal{R}_q \\ (B, \cdot) \in \text{sup}(\text{Gen})}} 1 - \Pr_{(s_1, s_2) \leftarrow \mathcal{D}_{\Lambda(B)}(c, 0), s} [\| (s_1, s_2) \|_2 > \beta] \cdot (1 + 2\epsilon) \quad (\text{Lemma 7, Lemma 10}) \\
&\geq \min_{\substack{c \in \mathcal{R}_q \\ (B, \cdot) \in \text{sup}(\text{Gen})}} 1 - \left(\frac{\rho_s(\Lambda(B))}{\rho_s(\Lambda(B) + t)} \cdot \left(\sqrt{e^{1-\tau^2}\tau^2} \right)^{2n} \cdot (1 + 2\epsilon) \right), \text{ for } t \in \Lambda(B)_{(c, 0)} \quad \left(\text{Lemma 2 and } \beta = \tau s \sqrt{2n} \right) \\
&\geq 1 - \left(\frac{1+\epsilon}{1-\epsilon} \cdot \left(\sqrt{e^{1-\tau^2}\tau^2} \right)^{2n} \cdot (1 + 2\epsilon) \right). \quad \left(\text{Lemma 5, } \det(\Lambda(B)) = \det(\Lambda(B) + t), \right. \\
&\quad \left. \text{and } s \geq \eta_\epsilon(\Lambda(B)) \right)
\end{aligned}$$

For FALCON⁺-512 and $\lambda = 128$, setting $\epsilon = (2^{64} \cdot 128)^{-1/2}$, $\tau = 1.1$, and $n = 512$ yields

$$p_{\text{PreSmp},\beta} \geq 1 - 2^{-14.31}.$$

For FALCON⁺-1024 and $\lambda = 256$, setting $\epsilon = (2^{64} \cdot 256)^{-1/2}$, $\tau = 1.1$, and $n = 1024$ analogously yields

$$p_{\text{PreSmp},\beta} \geq 1 - 2^{-28.63}.$$

When the following condition is satisfied:

$$Q_s \leq C_s p_{\text{PreSmp},\beta}, \quad (3)$$

Hoeffding's inequality can be applied to obtain a tail bound on the probability of observing at most Q_s successes in C_s independent Bernoulli trials. Specifically, the bound is given by,

$$\sum_{i=0}^{Q_s} \binom{C_s}{i} (1 - p_{\text{PreSmp},\beta})^{C_s-i} (p_{\text{PreSmp},\beta})^i \leq \exp \left(-2C_s \left(p_{\text{PreSmp},\beta} - \frac{Q_s}{C_s} \right)^2 \right) \quad (4)$$

where Q_s is the number of successes, C_s is the number of trials, and $p_{\text{PreSmp},\beta}$ is the probability of success in each trial. To satisfy the condition of Equation (3), C_s is set as follows,

$$C_s := 2^{64} + 2^{50} \geq \frac{2^{64}}{1 - 2^{-14.31}} \geq \frac{Q_s}{p_{\text{PreSmp},\beta}}.$$

Finally, the bound in Equation (4) is verified as follows,

$$\exp \left(-2 \cdot (2^{64} + 2^{50}) \left(1 - 2^{-14.31} - \frac{2^{64}}{2^{64} + 2^{50}} \right)^2 \right) \ll 2^\lambda \quad (\text{for } \lambda = 128).$$

Similarly, setting $C_s := 2^{64} + 2^{36}$ suffices when $p_{\text{PreSmp},\beta} \geq 1 - 2^{-28.63}$ and $\lambda = 256$. ■

E.2 Proof of Corollary 3

Corollary 3 (Rényi Loss for Falcon⁺-512 (Preimage Sampler) in Thm. 1). For $\epsilon \geq 2^{-\lambda} = 2^{-120}$, $r_p = R_{a_p}(\text{PreSmp} \parallel \mathcal{D})$, $C_s = 2^{64} + 2^{50}$, and the parameters for FALCON⁺-512, the Rényi argument for

$$r_p^{C_s} \epsilon^{\frac{a_p-1}{a_p}}$$

loses at most 3.5 bits for an order $a_p \approx 72.96$.

Proof. By Lemma 11 we need to solve

$$\min_{a_p > 1} C_s \cdot \log(R_{a_p}(\text{PreSmp} \parallel \mathcal{D})) + \frac{\lambda}{a_p}.$$

By Corollary 1 we can upper bound R_{a_p}

$$\min_{a_p > 1} C_s \cdot \log(1 + 2a_p\epsilon^2) + \frac{\lambda}{a_p}.$$

Differentiating with respect to a_p gives

$$\frac{2 \cdot C_s \cdot \epsilon^2}{\ln(2) \cdot (2a_p\epsilon^2 + 1)} - \frac{\lambda}{a_p^2}.$$

Setting the derivative to 0 and rearranging the terms yields

$$0 = 2a_p^2 C_s \epsilon^2 - \lambda \ln(2) - 2a_p \epsilon^2 \lambda \ln(2).$$

With the condition $a_p > 1$ the solution of the quadratic equation is

$$a_p = \frac{\lambda \epsilon^2 \ln(4) + \sqrt{8C_s \lambda \epsilon^2 \ln(2) + \lambda^2 \epsilon^4 \ln^2(4)}}{4C_s \epsilon^2} \quad (5)$$

Plugging $\lambda = 120$, $\epsilon = 1/\sqrt{2^{64} \cdot 128} = 2^{-35.5}$ and $C_s = 2^{64} + 2^{50}$ into Equation (5) gives

$$a_p \approx 72.96$$

and thus a bit loss of at most

$$C_s \cdot \log(1 + 2 \cdot 72.96 \cdot \epsilon^2) + \frac{120}{72.96} \leq 3.29.$$

■

E.3 Proof of Corollary 4

Corollary 4 (Rényi Loss for Falcon⁺-512 (Uniformity) in Thm. 1). For $\varepsilon \geq 2^{-\lambda} = 2^{-116.5}$, $r_u = R_{a_u}(\mathcal{U}(\mathcal{R}_q) \parallel \mathcal{U}_h)$, $C_s = 2^{64} + 2^{50}$, and the parameters for FALCON⁺-512, the Rényi argument for

$$r_u^{C_s} \varepsilon^{\frac{a_u - 1}{a_u}}$$

loses at most 3.5 bits for an order $a_u \approx 71.73$.

Proof. The corollary can be proved similar to the proof of Corollary 3 except that the Rényi divergence is upper bounded using Corollary 2. This leads to minimizing

$$\min_{a_u > 1} C_s \cdot \log\left(1 + \frac{2 \cdot a_u \cdot \epsilon^2}{(1 - \epsilon)^2}\right) + \frac{\lambda}{a_u},$$

which yields the statement. ■

E.4 Additional Rényi Corollaries

For the t - $\mathcal{R}$ -ISIS term we obtain a security of 278 bits, i.e. we can assume that the Rényi argument of the preimage sampler needs to preserve at most $\lambda = 278$ bits.

Corollary 5 (Rényi Loss for Falcon⁺-1024 (Preimage Sampler) in Thm. 1). For $\varepsilon \geq 2^{-\lambda} = 2^{-278}$, $r_p = R_{a_p}(\text{PreSmp} \parallel \mathcal{D})$, $C_s = 2^{64} + 2^{36}$, and the parameters for FALCON⁺-1024, the Rényi argument for

$$r_p^{C_s} \varepsilon^{\frac{a_p-1}{a_p}}$$

loses at most 4 bits for an order $a_p \approx 157.05$.

Proof. The proof works as the proof of Corollary 3 with different parameters. ■

Since we already lost 4 bits when unfolding the Rényi argument for the preimage sampler, we need to apply the following corollary with a security level of only 274 bits.

Corollary 6 (Rényi Loss for Falcon⁺-1024 (Uniformity) in Thm. 1). For $\varepsilon \geq 2^{-\lambda} = 2^{-274}$, $r_u = R_{a_u}(\mathcal{U}(\mathcal{R}_q) \parallel \mathcal{U}_h)$, $C_s = 2^{64} + 2^{36}$, and the parameters for FALCON⁺-1024, the Rényi argument for

$$r_u^{C_s} \varepsilon^{\frac{a_u-1}{a_u}}$$

loses at most 4 bits for an order $a_u \approx 155.92$.

Proof. The proof works as the proof of Corollary 4 with different parameters. ■

F Samplers

Here we recall the KLEIN SAMPLER [Kle00] and the FFO SAMPLER and prove similar results as in [Pre17]. However, we first present the Gram–Schmidt orthogonalization and LDL decomposition at the core of the samplers.

To present the FFO SAMPLER, we require some additional notation. The canonical embedding of $a \in \mathbb{R}[X]/(X^n + 1)$ is $\mathbf{a} = (a(\zeta))_{\zeta^n+1=0} \in \mathbb{C}^{n/2}$, since there are $n/2$ complex primitive elements ζ such $\zeta^n + 1 = 0$ when n is a power of two. The mapping between $\mathbb{R}[X]/(X^n + 1)$ and $\mathbb{C}^{n/2}$ is called the canonical embedding. With every element $a \in \mathbb{R}[X]/(X^n + 1)$, there is an adjoint $a^\dagger$, uniquely defined by the condition $a^\dagger(\zeta) = \overline{a(\zeta)}$ for all ζ such $\zeta^n + 1 = 0$. It allows to define an inner product over $\mathbb{C}^{n/2}$ $\langle \mathbf{a}, \mathbf{b} \rangle = \sum_{\zeta^n+1=0} a(\zeta) \overline{b(\zeta)}$ and the associated norm $\|\mathbf{a}\| = \sqrt{\langle \mathbf{a}, \mathbf{a} \rangle}$. Equipped with this scalar product, the embedding allows to view the ring of integers $\mathbb{Z}[X]/(X^n + 1)$ as a euclidean lattice in $\mathbb{C}^{n/2}$, or in $\mathbb{R}^n$. From now on, lowercase letters denote polynomials, and bold lowercase letters denote vectors as is customary in lattice literature.

F.1 Orthogonalizations

GRAM–SCHMIDT ORTHOGONALIZATION. For any linearly independent vectors $(\mathbf{b}_1, \dots, \mathbf{b}_n)$ there exist orthogonal vectors $(\tilde{\mathbf{b}}_1, \dots, \tilde{\mathbf{b}}_n)$ such that

$$\forall i \in [n]: \quad \text{span}(\mathbf{b}_1, \dots, \mathbf{b}_i) = \text{span}(\tilde{\mathbf{b}}_1, \dots, \tilde{\mathbf{b}}_i)$$

and one can compute such orthogonal vectors with the following formula:

$$\forall i \in [n]: \quad \tilde{\mathbf{b}}_i = \mathbf{b}_i - \sum_{j=1}^{i-1} \frac{\langle \mathbf{b}_i, \tilde{\mathbf{b}}_j \rangle_2}{\langle \tilde{\mathbf{b}}_j, \tilde{\mathbf{b}}_j \rangle_2} \tilde{\mathbf{b}}_j.$$

One can write this in matrix form: $\mathbf{B} = \tilde{\mathbf{B}} \cdot \mathbf{L}$, where $\mathbf{L}$ is unit upper triangular, and $\mathbf{B}$ and $\tilde{\mathbf{B}}$ are the column matrices of $(\mathbf{b}_1, \dots, \mathbf{b}_n)$ and $(\tilde{\mathbf{b}}_1, \dots, \tilde{\mathbf{b}}_n)$, respectively. The unit upper triangular condition on $\mathbf{L}$ makes the Gram–Schmidt decomposition unique.

LDL[†] DECOMPOSITION. For any matrix $\mathbf{A}$ such that $\mathbf{A} = \mathbf{A}^\dagger$ (Hermitian matrix), there exists a unique pair $(\mathbf{L}, \mathbf{D})$, where $\mathbf{L}$ is upper unit triangular and $\mathbf{D}$ is diagonal, such that $\mathbf{A} = \mathbf{L}^\dagger \mathbf{D} \mathbf{L}$. It is worth noting that the Gram–Schmidt decomposition and the LDL[†] decomposition are closely related. By the uniqueness of both decompositions, one can identify $\mathbf{L}$ in $\mathbf{B} = \tilde{\mathbf{B}} \cdot \mathbf{L}$ and in the LDL[†] decomposition of $\mathbf{A} = \mathbf{B}^\dagger \mathbf{B}$, which is an Hermitian matrix, as $\mathbf{A} = \mathbf{B}^\dagger \mathbf{B} = \mathbf{L}^\dagger \tilde{\mathbf{B}}^\dagger \tilde{\mathbf{B}} \mathbf{L} = \mathbf{L}^\dagger \mathbf{D} \mathbf{L}$ and $\tilde{\mathbf{B}}$ is an orthogonal basis.

F.2 Klein Sampler

Lemma 17 (Isotropic Gaussian in Orthogonal Basis). For any $s \in \mathbb{R}$, for any $\mathbf{t} \in \mathbb{R}^n$ and any orthogonal basis $\tilde{\mathbf{B}} = (\tilde{\mathbf{b}}_1, \dots, \tilde{\mathbf{b}}_n)$, where $z_i \tilde{\mathbf{b}}_i \leftarrow \mathcal{D}_{\mathbb{Z}\tilde{\mathbf{b}}_i, s, t_i \tilde{\mathbf{b}}_i}$ then $\mathbf{B}\mathbf{z} \leftarrow \mathcal{D}_{\Lambda(\tilde{\mathbf{B}}), s, \mathbf{t}\tilde{\mathbf{B}}}$.

Proof. We denote by $(Z_i)_{i \in [n]}$ the random choices made for each coordinate of $\mathbf{z}$. Then, by the definition of $\mathcal{D}_{\mathbb{Z}\tilde{\mathbf{b}}_i, s, t_i \tilde{\mathbf{b}}_i}$, we have $\Pr[Z_i = z_i] = \frac{\rho_s((z_i - t_i)\tilde{\mathbf{b}}_i)}{\rho_s(\mathbb{Z}\tilde{\mathbf{b}}_i)}$. Since each coordinate is sampled independently, it follows that $\Pr[\mathbf{Z} = \mathbf{z}] = \prod_{i \in [n]} \Pr[Z_i = z_i] = \frac{\prod_{i \in [n]} \rho_s((z_i - t_i)\tilde{\mathbf{b}}_i)}{\prod_{i \in [n]} \rho_s(\mathbb{Z}\tilde{\mathbf{b}}_i)}$. The numerator in the above expression simplifies as $\prod_{i \in [n]} \rho_s((z_i - t_i)\tilde{\mathbf{b}}_i) = \rho_s\left(\sum_{i \in [n]} (z_i - t_i)\tilde{\mathbf{b}}_i\right) = \rho_s((\mathbf{z} - \mathbf{t})\tilde{\mathbf{B}})$ where the penultimate equality holds since the vectors $\tilde{\mathbf{b}}_i$ are orthogonal and as Gaussian function ρ_s is additive for orthogonal vectors. Finally, by probability normalization, we obtain $\Pr[\mathbf{Z} = \mathbf{z}] = \frac{\rho_s((\mathbf{z} - \mathbf{t})\tilde{\mathbf{B}})}{\rho_s(\Lambda(\tilde{\mathbf{B}}))}$ which is exactly the expected distribution. ■

This sampler, studied in [Kle00, GPV08], is an adaptation of Babai’s nearest plane algorithm that introduces Gaussian sampling so that the output distribution does not reveal information about the secret basis used as the CVP trapdoor.

Algorithm 1 KLEIN_{B,s}($\mathbf{t} \in \mathbb{R}^n$)

Require: $s \geq \eta_\epsilon(\mathbb{Z}^n) \cdot \|\mathbf{B}\|_{GS}$, the Gram–Schmidt decomposition $\mathbf{B} = \tilde{\mathbf{B}} \cdot \mathbf{L}$ and the values $s_j = s / \|\tilde{\mathbf{b}}_j\|_2$ for $j \in [n]$
Ensure: A vector $\mathbf{z}$ such that $\mathbf{B}\mathbf{z} \leftarrow \mathcal{D}_{\Lambda(\mathbf{B}), s, \mathbf{B}\mathbf{t}}$

- 1: **for** $j \in \{n, \dots, 1\}$ **do**
 - 2: $t'_j \leftarrow t_j + \sum_{i > j} L_{ij}(t_i - z_i)$
 - 3: $z_j \leftarrow D_{\mathbb{Z}, s_j, t'_j}$
 - 4: **return** $\mathbf{z}$
-

For completeness and because the proof of the KLEIN SAMPLER is similar to the one of the FFO SAMPLER, we recall its output distribution and some intermediate lemmas.

Lemma 18 ([GPV08, Lem 4.4]). Let $\mathbf{B}$ be the lattice basis, and $\tilde{\mathbf{B}}$ its orthogonalisation basis. For any input $\mathbf{t} \in \mathbb{R}^n$, $\mathbf{t}' \in \mathbb{R}^n$ as defined in Algorithm 1, and any output $\mathbf{z} \in \mathbb{Z}^n$ of KLEIN_{L,s},

$$\mathbf{B} \cdot (\mathbf{z} - \mathbf{t}) = \tilde{\mathbf{B}} \cdot (\mathbf{z} - \mathbf{t}').$$

Proof. Using the transition matrix $\mathbf{L}$, one can compute the coordinates of $\mathbf{B} \cdot (\mathbf{z} - \mathbf{t})$ in basis $\tilde{\mathbf{B}}$:

$$\begin{aligned}
(\mathbf{L} \cdot (\mathbf{z} - \mathbf{t}))_i &= \sum_{j=1}^n L_{ij} \cdot (\mathbf{z} - \mathbf{t})_j \\
&= z_i - t_i + \sum_{j>i} L_{ij} \cdot (\mathbf{z} - \mathbf{t})_j \quad (\mathbf{L} \text{ is unit upper triangular}) \\
&= z_i - t'_i \\
&= (\mathbf{z} - \mathbf{t}')_i.
\end{aligned}$$

Therefore,

$$\mathbf{B} \cdot (\mathbf{z} - \mathbf{t}) = \tilde{\mathbf{B}} \cdot \mathbf{L} \cdot (\mathbf{z} - \mathbf{t}) = \tilde{\mathbf{B}} \cdot (\mathbf{z} - \mathbf{t}').$$

This completes the proof. ■

Lemma 19 ([GPV08, Lem 4.5]). For any input $\mathbf{t} \in \mathbb{R}^n$, any $\mathbf{z} \in \mathbb{Z}^n$, and any $s \geq \eta_\epsilon(\mathbb{Z}^n) \cdot \|\mathbf{B}\|_{GS}$, with $\epsilon \in (0, 1/4)$, and Gram-Schmidt decomposition $\mathbf{B} = \tilde{\mathbf{B}} \cdot \mathbf{L}$ and the values $s_j = s / \|\tilde{\mathbf{b}}_j\|_2$ for $j \in [n]$, the probability that $\text{KLEIN}_{\mathbf{B},s}$ outputs $\mathbf{z} \in \mathbb{R}^n$ is exactly

$$\rho_s(\mathbf{B} \cdot (\mathbf{t} - \mathbf{z})) \cdot \prod_{i \in [n]} \frac{1}{\rho_{s_i, t'_i}(\mathbb{Z})},$$

where the values s_i, t'_i are as in the execution of $\text{KLEIN}_{\mathbf{B},s}(\mathbf{t}) \rightarrow \mathbf{z}$.

Proof. Consider the event E that $\text{KLEIN}_{\mathbf{L},s}(\mathbf{t})$ is exactly $\mathbf{z}$. We denote by $(Z_i)_{i \in [n]}$ the random choices made by $\text{KLEIN}_{\mathbf{L},s}(\mathbf{t})$, the event E is exactly the event where each $Z_i = z_i$ for $i \in [n]$. Now for each i , the probability that $Z_i = z_i$ conditioned on $Z_j = z_j$ for each $j = n, \dots, i+1$, is exactly $\mathcal{D}_{\mathbb{Z}, s_i, t'_i}(z_i)$. Therefore the probability of E is

$$\prod_{i \in [n]} \mathcal{D}_{\mathbb{Z}, s_i, t'_i}(z_i) = \frac{\prod_{i \in [n]} \rho_{s_i, t'_i}(z_i)}{\prod_{i \in [n]} \rho_{s_i, t'_i}(\mathbb{Z})}.$$

The numerator in the above expression is

$$\begin{aligned}
\prod_{i \in [n]} \rho_{s_i, t'_i}(z_i) &= \prod_{i \in [n]} \rho_s((t'_i - z_i) \cdot \|\tilde{\mathbf{b}}_i\|) \\
&= \rho_s \left(\sum_{i \in [n]} \tilde{\mathbf{b}}_i \cdot (t'_i - z_i) \right) \quad (\text{orthogonality of } \tilde{\mathbf{B}}) \\
&= \rho_s(\tilde{\mathbf{B}} \cdot (\mathbf{t}' - \mathbf{z})) \\
&= \rho_s(\mathbf{B} \cdot (\mathbf{t} - \mathbf{z})). \quad (\text{Lemma 18})
\end{aligned}$$

This completes the proof. ■

Lemma 20 (Rényi Divergence of Klein Sampler [Pre17, Lem. 6]). Let n be a positive integer, $a > 1$, and $\epsilon \in (0, 1/4)$. Then for the $\text{KLEIN SAMPLER PreSmp}$ and the lattice $\mathbf{\Lambda} = \mathbf{\Lambda}(\mathbf{B})_{(\mathbf{c}, \mathbf{0})}$, for any basis $\mathbf{B} \in \mathbb{Z}^{2n \times 2n}$, standard deviation $s \geq \eta_\epsilon(\mathbb{Z}^{2n}) \cdot \|\mathbf{B}\|_{GS}$, and arbitrary syndrome $\mathbf{c} \in \mathcal{R}_q$, the Rényi divergence is bounded by

$$R_a(\text{PreSmp}(\mathbf{B}, s, (\mathbf{c}, \mathbf{0})) \parallel \mathcal{D}_{\mathbf{\Lambda}, s}) \lesssim 1 + 2a\epsilon^2.$$

Proof. In the proof, $\mathbf{t} = (\mathbf{c}, \mathbf{0})$. As stated in previous lemma, the probability that $\text{KLEIN}_{\mathbf{L},s}(\mathbf{t})$ outputs a given $\mathbf{z}$ is exactly

$$\rho_s(\mathbf{B} \cdot (\mathbf{t} - \mathbf{z})) \cdot \prod_{i \in [2n]} \frac{1}{\rho_{s_i, t'_i}(\mathbb{Z})}.$$

As $s_i = s/\|\tilde{\mathbf{b}}_i\|_2 \geq s/\|\mathbf{B}\|_{GS}$, by assumption $s_i \geq \eta_\epsilon(\mathbb{Z}^{2n}) \geq \eta_{\epsilon/2n}(\mathbb{Z})$, therefore $\rho_{s_i, t'_i}(\mathbb{Z}) \in \left[\frac{1-\epsilon/2n}{1+\epsilon/2n}, 1\right] \cdot \rho_{s_i}(\mathbb{Z})$ by [MR04, Lemma 4.4]. Since $\mathcal{D}_{\Lambda,s}(\mathbf{B} \cdot \mathbf{z})$ is proportional to $\rho_s(\mathbf{B} \cdot (\mathbf{t} - \mathbf{z}))$ and as $s_i = s/\|\tilde{\mathbf{b}}_i\|_2 \geq s/\|\mathbf{B}\|_{GS}$, by assumption $s_i \geq \eta_\epsilon(\mathbb{Z}^{2n}) \geq \eta_{\epsilon/2n}(\mathbb{Z})$, therefore $\rho_{s_i, t'_i}(\mathbb{Z}) \in \left[\frac{1-\epsilon/2n}{1+\epsilon/2n}, 1\right] \cdot \rho_{s_i}(\mathbb{Z})$ by [MR04, Lemma 4.4]. Since $\mathcal{D}_{\Lambda,s}(\mathbf{B} \cdot \mathbf{z})$ is proportional to $\rho_s(\mathbf{B} \cdot (\mathbf{t} - \mathbf{z}))$ and both $\mathcal{D}_{\Lambda,s}$ and $\text{PreSmp}(\mathbf{B}, s, (\mathbf{c}, \mathbf{0}))$ sum up to one, we have

$$\left(\frac{1-\epsilon/2n}{1+\epsilon/2n}\right)^{2n} \leq \frac{\text{PreSmp}(\mathbf{B}, s, (\mathbf{c}, \mathbf{0}))}{\mathcal{D}_{\Lambda,s}} \leq \left(\frac{1+\epsilon/2n}{1-\epsilon/2n}\right)^{2n}.$$

Therefore,

$$\delta_{RE}(\text{PreSmp}(\mathbf{B}, s, (\mathbf{c}, \mathbf{0})), \mathcal{D}_{\Lambda,s}) \leq \left(\frac{1+\epsilon/2n}{1-\epsilon/2n}\right)^{2n} - 1 \approx 2\epsilon,$$

from which we can conclude the proof by using the relative error lemma. $\blacksquare$

F.3 FFO Sampler

The main drawback of the KLEIN SAMPLER is its quadratic complexity. The FFO SAMPLER algorithm is an adaptation of the KLEIN SAMPLER that exploits the structure of the NTRU matrices using an advanced Fast Fourier Transform depicted in [DP16]. We first present how the KLEIN SAMPLER is modified.

The KLEIN SAMPLER uses the orthogonalization of the matrix $\mathbf{B}$ to sample $2n$ vectors in the $2n$ orthogonal lattices $\Lambda(\tilde{\mathbf{b}}_j)$ (this is what the line $z_j \leftarrow \mathcal{D}_{\mathbb{Z}, s_j, t'_j}$ does). One can take advantage of the structure of NTRU lattices and modify this procedure to achieve quasi-linear complexity. Given a basis

$$\mathbf{B} = \begin{pmatrix} g & G \\ -f & -F \end{pmatrix} \in \mathcal{R}_q^{2 \times 2},$$

one wants to sample a vector from the lattice $\Lambda(\mathcal{A}(\mathbf{B}))$. To do this, one can apply Klein's algorithm in dimension 2 on the basis $\mathbf{B}$. First, one compute a block orthogonalisation of the NTRU basis $\mathbf{B}$: $\mathbf{B} = (\tilde{\mathbf{b}}_1 \tilde{\mathbf{b}}_2) \cdot \mathbf{L}$. Then, it is sufficient to be able to sample in the lattices $\Lambda(\mathcal{A}(\tilde{\mathbf{b}}_1))$ and $\Lambda(\mathcal{A}(\tilde{\mathbf{b}}_2))$ using two recursive calls, which would yield our quasi-linear algorithm.

The obstacle is that $\mathcal{A}(\tilde{\mathbf{b}}_1)$ and $\mathcal{A}(\tilde{\mathbf{b}}_2)$ cannot be interpreted as matrices of the image of $\mathcal{A}$, as the coefficients of $\tilde{\mathbf{b}}_i$ are not integers. To overcome this issue, new operators analogous to c (coefficient embedding of an element of $\mathcal{R}_q$) and $\mathcal{A}$ (anticirculant matrix) are introduced: $V_{d/d'}$ and $M_{d/d'}$, where d and d' are powers of 2. It is possible to define a more general V but for NTRU, powers of 2 are sufficient. Informally, V can be viewed as iterations of the “split” operator used in the fast Fourier transform.

To go down the tower rings, let us introduce the notation $\mathcal{R}_q(d)$ to denote the subring of dimension 2^d of $\mathcal{R}_q(k)$ with $n = 2^k$. The tower-of-fields of the cyclotomic field $\mathbb{Q}[X]/(X^n + 1)$ for $n = 2^k$ has corresponding subfields that we will write as $Q(d)$. This notation is the same as in the forthcoming FN-DSA standard.

The following definitions and properties come from [DP16].

- Define $V_{2^d/2^{d-1}}$, for $a \in \mathcal{R}_q(d)$, $V_{2^d/2^{d-1}}(a) = \begin{pmatrix} p_0 \\ p_1 \end{pmatrix}$ where $p_0, p_1 \in \mathcal{R}_q(d-1)$ are the unique polynomials such that $a(X) = p_0(X^2) + X p_1(X^2)$.
- Define $V_{2^d/2^{d'}}$ with $d \geq d'$ recursively as the identity if $d = d'$, otherwise for $a \in \mathcal{R}_q(d)$, $V_{2^d/2^{d'}}(a)$ is the result of the coefficient-wise application of $V_{2^{d-1}/2^{d'}}(a)$ to $V_{2^d/2^{d-1}}(a)$.

- For $a \in \mathcal{R}_q(d)$, if $V_{2^d/2^{d-1}}(a) = \begin{pmatrix} p_0 \\ p_1 \end{pmatrix}$ then

$$M_{2^d/2^{d-1}}(a) = \begin{bmatrix} p_0 & X p_1 \\ p_1 & p_0 \end{bmatrix} = [V_{2^d/2^{d-1}}(a) \ V_{2^d/2^{d-1}}(Xa)] \in \mathcal{R}_q(d-1)^{2 \times 2}.$$

- As for V , define $M_{2^d/2^{d'}}$ with $d \geq d'$ recursively as the identity if $d = d'$, otherwise for $a \in \mathcal{R}_q(d)$, $M_{2^d/2^{d'}}(a)$ is the result of the coefficient wise application of $M_{2^{d-1}/2^{d'}}$ to $M_{2^d/2^{d-1}}(a)$.
- V and M generalize to vectors and matrices in a coefficient-wise manner.

Observe that, as stated in [DP16], the following properties hold:

- $M(\mathbf{A} \cdot \mathbf{B}) = M(\mathbf{A}) \cdot M(\mathbf{B})$
- $V(ab) = M(a) \cdot V(b)$
- V is an isometry: $\langle V(\mathbf{a}), V(\mathbf{b}) \rangle_2 = \langle \mathbf{a}, \mathbf{b} \rangle_2$

In the FFO SAMPLER algorithm, instead of sampling in the lattice generated by $\mathcal{A}(\mathbf{B})$, one samples in the lattice generated by $M_{2^k/1}(\mathbf{B})$. The operator $M_{2^k/1}$ enables the recursive call discussed earlier. When sampling in the lattice generated by $M_{2^k/1}(\tilde{\mathbf{b}}_i)$, one can decompose $\tilde{\mathbf{b}}_i \in \mathcal{R}_q(k)^m$ as a matrix $\tilde{\mathbf{B}}_i \in \mathcal{R}_q(k-1)^{2 \times 2m}$ using the partial operator $M_{2^k/2^{k-1}}$. By definition of M , the lattices generated by $M_{2^k/1}(\tilde{\mathbf{b}}_i)$ and $M_{2^{k-1}/1}(\tilde{\mathbf{B}}_i)$ are the same.

COMPACT $LDL^\dagger$ DECOMPOSITION. These linearization operators provide a compact way to express the $LDL^\dagger$ decomposition. This decomposition is given by algorithm 2:

Algorithm 2 $\text{ffLDL}^\dagger(Q)$

Require: A positive-definite self-adjoint matrix $Q \in \mathcal{Q}(k)^{2 \times 2}$

Ensure: A binary tree $\mathcal{T}$.

- 1: $(L_{10}, D_{00}, D_{11}) \leftarrow \text{LDL}^\dagger(Q)$
 - 2: $\mathcal{T}.\text{value} \leftarrow L_{10}$
 - 3: **if** $d = 1$ **then**
 - 4: $\mathcal{T}.\text{leftchild} \leftarrow D_{00}$
 - 5: $\mathcal{T}.\text{rightchild} \leftarrow D_{11}$
 - 6: **return** $\mathcal{T}$
 - 7: **else**
 - 8: $Q_0 \leftarrow M_{2^d/2^{d-1}}(D_{00})$
 - 9: $Q_1 \leftarrow M_{2^d/2^{d-1}}(D_{11})$
 - 10: $\mathcal{T}.\text{leftchild} \leftarrow \text{ffLDL}^\dagger(Q_0)$
 - 11: $\mathcal{T}.\text{rightchild} \leftarrow \text{ffLDL}^\dagger(Q_1)$
 - 12: **return** $\mathcal{T}$
-

This algorithm computes a “compact LDL decomposition.” Indeed, consider $\mathbf{B} \in \mathcal{R}_q(k)^{2 \times m}$ by writing

$\text{ffLDL}^\dagger(\mathbf{B}^\dagger \cdot \mathbf{B})$ in the form $\begin{matrix} & & L \\ & \swarrow & \searrow \\ L_0 & & L_1 \\ \swarrow \quad \searrow & & \swarrow \quad \searrow \\ L_{00} & L_{01} & L_{10} & L_{11} \\ \vdots & \vdots & \vdots & \vdots \\ D_1 & \dots & \dots & D_{2n} \end{matrix}$, and with the properties of $\text{LDL}^\dagger$, an immediate induction

provides the following results:

$$M_{2^d/2^{d-i}}(\mathbf{B}) = \tilde{\mathbf{B}}_i \cdot \mathbf{L}_i$$

where $\tilde{\mathbf{B}}_i$ is the Gram–Schmidt matrix of $M_{2^d/2^{d-i}}(\mathbf{B})$, and $\mathbf{L}_i$ is defined recursively by

$$\mathbf{L}_0 = \begin{bmatrix} 1 & L \\ 0 & 1 \end{bmatrix}, \quad \forall i > 0: \mathbf{L}_i = \begin{bmatrix} 1 & \underbrace{L_0 \dots 0}_i & & & 0 \\ 0 & 1 & & & \\ & & 1 & L_0 \dots 01 & \\ & & 0 & 1 & \\ & & & & \ddots \\ 0 & & & & 1 & L_1 \dots 1 \\ & & & & 0 & 1 \end{bmatrix} \cdot M_{2^{k-i+1}/2^{k-i}}(\mathbf{L}_{i-1})$$

Moreover, as $\text{Diag}(D_1, \dots, D_{2n}) = \tilde{\mathbf{B}}_n^\dagger \tilde{\mathbf{B}}_n$ the $(D_1, \dots, D_{2n})$ are the squares of the norms of the Gram–Schmidt basis vectors of $M_{2^k/1}(\mathbf{B})$.

FFO SAMPLER. We can now put into practice the suggested idea to efficiently sample vectors from a lattice generated by the secret basis $\mathbf{B}$.

Algorithm 3 ffSampling($\mathbf{t}, \mathcal{T}$)

Require: An element $\mathbf{t} = [t_0, t_1] \in Q(k)^2$, and an LDL tree $\mathcal{T}$.

Ensure: An element $\mathbf{z} = [z_0, z_1] \in \mathcal{R}_q(k)^2$.

```

1: if  $k = 1$  then
2:    $l \leftarrow \mathcal{T}.\text{value}$ 
3:    $s_0 \leftarrow s / \sqrt{\mathcal{T}.\text{leftchild}}$ 
4:    $s_1 \leftarrow s / \sqrt{\mathcal{T}.\text{rightchild}}$ 
5:    $z_1 \leftarrow D_{\mathbb{Z}, t_1, s_1}$ 
6:    $t'_0 \leftarrow t_0 + l \cdot (t_1 - z_1)$ 
7:    $z_0 \leftarrow D_{\mathbb{Z}, t'_0, s_0}$ 
8:   return ( $\mathbf{z} = [z_0, z_1]$ )
9: ( $L, \mathcal{T}_0, \mathcal{T}_1$ )  $\leftarrow (\mathcal{T}.\text{value}, \mathcal{T}.\text{leftchild}, \mathcal{T}.\text{rightchild})$ 
10:  $\mathbf{t}_1 \leftarrow V_{k/\frac{k}{2}}(t_1)$ 
11:  $\mathbf{z}_1 \leftarrow \text{ffSampling}(\mathbf{t}_1, \mathcal{T}_1)$ 
12:  $\mathbf{z}_1 \leftarrow V_{k/\frac{k}{2}}^{-1}(\mathbf{z}_1)$ 
13:  $t'_0 \leftarrow t_0 + L \cdot (t_1 - z_1)$ 
14:  $\mathbf{t}'_0 \leftarrow V_{k/\frac{k}{2}}(t'_0)$ 
15:  $\mathbf{z}_0 \leftarrow \text{ffSampling}(\mathbf{t}'_0, \mathcal{T}_0)$ 
16:  $\mathbf{z}_0 \leftarrow V_{k/\frac{k}{2}}^{-1}(\mathbf{z}_0)$ 
17: return ( $\mathbf{z} = [z_0, z_1]$ )

```

Now, we can prove that the distribution of this new sampler has the same properties as the KLEIN SAMPLER with respect to the Rényi divergence.

Lemma 7 (Relative Error of FFO Sampler). Let n be a positive integer and $\epsilon \in (0, 1/4)$. Then the *relative error* of $\text{PreSmp} := \text{FFO SAMPLER}$ and the lattice $\mathbf{\Lambda} = \mathbf{\Lambda}(\mathbf{B})_{(\mathbf{c}, \mathbf{0})}$ for any basis $\mathbf{B} \in \mathbb{Z}^{2n \times 2n}$, standard deviation $s \geq \eta_\epsilon(\mathbb{Z}^{2n}) \cdot \|\mathbf{B}\|_{GS}$, and arbitrary syndrome $\mathbf{c} \in \mathcal{R}_q$ is bounded by

$$\delta_{RE}(\text{PreSmp}(\mathbf{B}, s, (\mathbf{c}, \mathbf{0})), \mathcal{D}_{\mathbf{\Lambda}, s}) \leq \left(\frac{1 + \epsilon/2n}{1 - \epsilon/2n} \right)^{2n} - 1 \approx 2\epsilon.$$

In order to prove this lemma, we show the same intermediate lemma as for the KLEIN SAMPLER (Lemma 19) by computing the distribution of the FFO SAMPLER. Consequently, since the last part of the proof is exactly the same as in Corollary 1, the theorem is proved. The vector $\mathbf{t}$ denotes $(\mathbf{c}, \mathbf{0})$.

Lemma 21. Let $\mathbf{B} = (\mathbf{b}_1, \mathbf{b}_2) \in Q(k)^{2 \times m}$ and $\tilde{\mathbf{B}} = (\tilde{\mathbf{b}}_1, \tilde{\mathbf{b}}_2)$ be its Gram-Schmidt orthogonalization. The vectors $\mathbf{z}$ and $\mathbf{t}' = \begin{pmatrix} t'_0 \\ t'_1 \end{pmatrix}$ in the first step of $\text{ffSampling}(\mathbf{t}, \mathcal{T}_{\mathbf{B}})$ satisfy

$$\mathbf{B} \cdot (\mathbf{z} - \mathbf{t}) = \tilde{\mathbf{B}} \cdot (\mathbf{z} - \mathbf{t}').$$

Proof. The proof is exactly the same as for Lemma 18, as each individual step of ffSampling is basically the KLEIN SAMPLER, where the calls to the Gaussian sampler over the integers are replaced by recursive calls. $\blacksquare$

Lemma 22. For any basis $\mathbf{B} \in Q(k)^{2 \times m}$, vector $\mathbf{t} \in Q(k)^2$, for any $\mathbf{z} = (\hat{z}_0, \hat{z}_1) \in R(k)^2$, and for any $s \geq \eta_\epsilon(\mathbb{Z}^n) \cdot \|\mathbf{B}\|_{GS}$, with the Gram-Schmidt decomposition $\mathbf{B} = \tilde{\mathbf{B}} \cdot \mathbf{L}$, the probability that $\text{ffSampling}(\mathbf{t}, \mathcal{T}_{\mathbf{B}})$ outputs $\mathbf{z}$ is exactly

$$\rho_s(\mathbf{B} \cdot (\mathbf{t} - \mathbf{z})) \cdot \prod_{i \in [2n]} \frac{1}{\rho_{\hat{s}_i, \hat{t}_i}(\mathbb{Z})},$$

where

- $\hat{s}_i, \hat{t}_i$ are the parameters of the distributions $D_{\mathbb{Z}, \hat{t}_i, \hat{s}_i}$ from which the execution of $\text{ffSampling} \rightarrow \mathbf{z}$ samples, in reverse order.
- $\mathcal{T}_{\mathbf{B}} = \text{ffLDL}^\dagger(\mathbf{B}\mathbf{B}^\dagger)$.

Proof. In this proof, will write $V = V_{k/\frac{k}{2}}$ and $M = M_{k/\frac{k}{2}}$ for conciseness. Consider the event E : the output of $\text{ffSampling}(\mathbf{t}, \mathcal{T}_{\mathbf{B}})$ is exactly (z_0, z_1) . We will prove the result inductively:

If $k = 1$, as the base case of ffSampling is exactly an execution of KLEIN of dimension 2, Lemma 19 concludes.

If $k \neq 1$, we denote by $\mathbf{Z}_0, \mathbf{Z}_1$ the random choices made by $\text{ffSampling}(\mathbf{t}, \mathcal{T}_{\mathbf{B}})$, E occurs if and only if $\mathbf{Z}_0 = z_0$ and $\mathbf{Z}_1 = z_1$. We note that

$$\begin{aligned} \Pr[\mathbf{Z}_1 = z_1] &= \Pr[V^{-1}(\text{ffSampling}(\mathbf{t}_1, \mathcal{T}_1)) = z_1], \text{ and} \\ \Pr[\mathbf{Z}_0 = z_0 \mid \mathbf{Z}_1 = z_1] &= \Pr[V^{-1}(\text{ffSampling}(\mathbf{t}'_0, \mathcal{T}_0)) = z_0]. \end{aligned}$$

Then using the bijectivity of V , we obtain

$$\Pr[E] = \Pr[\text{ffSampling}(\mathbf{t}_1, \mathcal{T}_1) = V(z_1)] \cdot \Pr[\text{ffSampling}(\mathbf{t}'_0, \mathcal{T}_0) = V(z_0)].$$

Now we would like to use the induction hypothesis. Looking at the algorithm $\text{ffLDL}^\dagger$, we know that $\mathcal{T}_0$ and $\mathcal{T}_1$ are the results of $\text{ffLDL}^\dagger(\tilde{\mathbf{B}}_0^\dagger \tilde{\mathbf{B}}_0)$ and $\text{ffLDL}^\dagger(\tilde{\mathbf{B}}_1^\dagger \tilde{\mathbf{B}}_1)$ with $\tilde{\mathbf{B}}_0 = M(\tilde{\mathbf{b}}_0)$ and $\tilde{\mathbf{B}}_1 = M(\tilde{\mathbf{b}}_1)$ where $(\tilde{\mathbf{b}}_0, \tilde{\mathbf{b}}_1)$ is the orthogonalization of $\mathbf{B}$. This allows to use induction hypothesis and write

$$\Pr[E] = \rho_s(\tilde{\mathbf{B}}_0 \cdot (\mathbf{t}_0 - V(z_0))) \cdot \rho_s(\tilde{\mathbf{B}}_1 \cdot (\mathbf{t}'_1 - V(z_1))) \cdot \prod_{i \in [2n]} \frac{1}{\rho_{\hat{s}_i, \hat{t}_i}(\mathbb{Z})}.$$

The numerator in the above expression is

$$\begin{aligned} & \rho_s(\tilde{\mathbf{B}}_0 \cdot (\mathbf{t}_0 - V(z_0))) \cdot \rho_s(\tilde{\mathbf{B}}_1 \cdot (\mathbf{t}'_1 - V(z_1))) \\ &= \rho_s(M(\tilde{\mathbf{b}}_0) \cdot (V(\mathbf{t}_0) - V(z_0))) \cdot \rho_s(M(\tilde{\mathbf{b}}_1) \cdot (V(\mathbf{t}'_1) - V(z_1))) && \text{(definition of } \tilde{\mathbf{B}}_0, \tilde{\mathbf{B}}_1, \mathbf{t}_0 \text{ and } \mathbf{t}'_1) \\ &= \rho_s(V((\mathbf{t}_0 - z_0) \cdot \tilde{\mathbf{b}}_0)) \cdot \rho_s(V((\mathbf{t}'_1 - z_1) \cdot \tilde{\mathbf{b}}_1)) && \text{(linearity of } V \text{ and } M) \\ &= \rho_s((\mathbf{t}_0 - z_0) \cdot \tilde{\mathbf{b}}_0) \cdot \rho_s((\mathbf{t}'_1 - z_1) \cdot \tilde{\mathbf{b}}_1) && (V \text{ isometry}) \\ &= \rho_s((\mathbf{t}_0 - z_0) \cdot \tilde{\mathbf{b}}_0 + (\mathbf{t}'_1 - z_1) \cdot \tilde{\mathbf{b}}_1) && \text{(orthogonality of } \tilde{\mathbf{b}}_0 \text{ and } \tilde{\mathbf{b}}_1 \text{ (Lemma 17))} \\ &= \rho_s((\mathbf{t}_0 - z_0) \cdot \mathbf{b}_0 + (\mathbf{t}'_1 - z_1) \cdot \mathbf{b}_1) && \text{(Lemma 21)} \\ &= \rho_s(\mathbf{B} \cdot (\mathbf{t} - \mathbf{z})). \end{aligned}$$

The orthogonality of $\tilde{\mathbf{B}}_0 = M(\tilde{\mathbf{b}}_0)$ and $\tilde{\mathbf{B}}_1 = M(\tilde{\mathbf{b}}_1)$ is inherited from orthogonality of $\tilde{\mathbf{b}}_0$ and $\tilde{\mathbf{b}}_1$. Indeed, as $\tilde{\mathbf{b}}_0 \cdot \tilde{\mathbf{b}}_1^\dagger = 0$ and M satisfies $M(A \cdot B) = M(A) \cdot M(B)$, it holds that $M(\tilde{\mathbf{b}}_0) \cdot M(\tilde{\mathbf{b}}_1)^\dagger = 0$. $\blacksquare$

Now, we can go back the proof of Lemma 7. As in the proof of Lemma 20, since the intermediate lemmas are similar, we get

$$\delta_{RE}(\text{PreSmp}(\mathbf{B}, s, (\mathbf{c}, \mathbf{0})), \mathcal{D}_{\mathbf{\Lambda}, s}) \leq \left(\frac{1 + \epsilon/2n}{1 - \epsilon/2n} \right)^{2n} - 1 \approx 2\epsilon.$$