# A Closer Look at Falcon

Phillip Gajland[1,2], Jonas Janneck[2], and Eike Kiltz[2]

[1] Max Planck Institute for Security and Privacy
[2] Ruhr University Bochum

26th February 2025

**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 lattice-based schemes.

Its security hinges on a Rényi divergence-based argument for Gaussian samplers, a core element of the scheme. 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.

This work takes a closer look at Falcon and demonstrates that introducing a few minor, conservative modifications allows for the first formal proof of the scheme in the random oracle model. 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.

Unfortunately, our analysis shows that despite our modification of Falcon-512 and Falcon-1024 we do not achieve *strong unforgeability* for either scheme. For *plain unforgeability* we are able to show that our modifications to Falcon-512 barely satisfy the claimed 120-bit security target and for Falcon-1024 we confirm the claimed security level. As such we recommend revisiting Falcon and its parameters.

# Contents

# 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 $y = \mathsf{H}(m)$ in the range of $f$, then outputs the signature $\sigma = f^{-1}(y)$. Verification consists of checking that $f(\sigma) = \mathsf{H}(m)$. 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 $\sigma$ are sampled from $f^{-1}(\mathsf{H}(m))$. By leveraging NTRU lattices, introduced by Hoffstein, Pipher, and Silverman [HPS98,HHP+03], Falcon benefits from their ring structure, allowing a reduction in public keys by a factor of $\mathcal{O}(n)$ and accelerating many computations by a factor of $\mathcal{O}(n/\log n)$. 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. 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 4. 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 signing, 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.[3] 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].

```
$\mathsf{Sgn}(sk, m)$
Sample salt $r$
repeat
    $s \xleftarrow{\$} f^{-1}(\mathsf{H}(r, m))$
until $\|s\|_2 \leq \beta$
$\sigma := (r, s)$
```

```
$\mathsf{Sgn}^+(sk, m)$
repeat
    Sample salt $r$
    $s \xleftarrow{\$} f^{-1}(\mathsf{H}(pk, r, m))$
until $\|s\|_2 \leq \beta$
$\sigma := (r, s)$
```

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

Security analysis and recommendations. We provide a thorough security analysis of Falcon$^+$ in the random oracle model. Using our tools, we obtain concrete security bounds derived from our theorems focusing on minimising the bit security loss due to Rényi divergence arguments. Our findings show that using our Rényi divergence tools combined with existing techniques, neither Falcon$^+$-512 (NIST Level 1) nor Falcon$^+$-1024 (NIST Level 5) provide *strong unforgeability*. For comparison, schemes like Dilithium [LDK+22] already meet strong unforgeability. In the case of *plain unforgeability*, Falcon$^+$-512 achieves only 113 bits of provable security. By reducing the number of allowed signing queries from $2^{64}$ to $2^{58}$, we increase this to 119 bits, nearing the claimed security level. For Falcon$^+$-1024, we prove that it meets 256 bits of security for *plain unforgeability*. An overview of the provable bit security is shown in Table 1.

Note that we do not present concrete attacks against Falcon, nor do we claim that a better security proof is impossible. Rather, we show that our tools combined with currently known proof techniques are insufficient to fully justify the target security claims of Falcon and Falcon$^+$.

| **Scheme** | **UF-CMA** | **SUF-CMA** |
|---|---|---|
| Falcon$^+$-512 ($Q_s = 2^{64}$) | 113 | 89 |
| Falcon$^+$-512 ($Q_s = 2^{58}$) | 119 | 94 |
| Falcon$^+$-1024 ($Q_s = 2^{64}$) | 256 | 0 |

**Table 1.** Provable bit security levels of Falcon$^+$-512 and Falcon$^+$-1024.

[3] 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.

### 1.2 Technical Overview

The Gentry-Peikert-Vaikuntanathan Framework. The GPV framework [GPV08] provides a method for constructing secure lattice-based signature schemes via the full-domain-hash (FDH) paradigm [BR96], commonly known as *"hash-and-sign"*. Central to this framework is a *"preimage sampleable trapdoor function"* defined as $f_{\boldsymbol{A}}(s) := \boldsymbol{A}s \mod q$ where $\boldsymbol{A} \in \mathbb{Z}_q^{n \times m}$. Here, each signature essentially corresponds to a short preimage of the hash of a message. More specifically, the public key $pk$ is a full-rank matrix $\boldsymbol{A} \in \mathbb{Z}_q^{n \times m}$ (with $n \leq m$) which defines a $q$-ary lattice $\boldsymbol{\Lambda}$. The secret key (or trapdoor) $sk$ is a matrix $\boldsymbol{B} \in \mathbb{Z}_q^{m \times m}$ generating the lattice orthogonal to $\boldsymbol{\Lambda} \mod q$, enabling the efficient inversion of $f_{\boldsymbol{A}}$. A signature on a message $m$ is a short Gaussian vector $s \in \mathbb{Z}_q^m$ such that $\mathsf{H}(m) = \boldsymbol{A}s \mod q$, where $\mathsf{H} : \{0,1\}^\star \to \mathbb{Z}_q^n$ is a hash function. Verification involves checking both the shortness of $s$ and that $f_{\boldsymbol{A}}(s) = \mathsf{H}(m)$.

The GPV Proof. 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) Short Integer Solution (SIS) assumption [Ajt96]. Security can be establish in two ways: (1) via *collision resistance*, reducing to SIS, or (2) via *one-wayness*, reducing to ISIS. The original work [GPV08] presented a *tight* proof of *strong unforgeability* for FDH, leveraging collision resistance.

Suppose, for the sake of contradiction, that an adversary $\mathsf{A}$ breaks the *strong unforgeability* of the signature scheme producing a forgery $(m^\star, s^\star)$ where $s^\star$ is short and $\mathsf{H}(m^\star) = \boldsymbol{A}s^\star \mod q$. We construct an adversary $\mathsf{B}$ that breaks SIS by finding a collision in $f_{\boldsymbol{A}}(s)$. Given a SIS instance $\boldsymbol{A} \in \mathbb{Z}_q^{n \times m}$, $\mathsf{B}$ runs adversary $\mathsf{A}$ on the public key $pk = \boldsymbol{A}$ and simulates the random oracle $\mathsf{H}$ and signing oracle as follows.

- The random oracle is programmed as follows: for each fresh query to $\mathsf{H}(m)$, $\mathsf{B}$ samples a Gaussian $s_m$ and returns $\mathsf{H}(m) := \boldsymbol{A}s_m \mod q$ to $\mathsf{A}$. Crucially, by the *"leftover hash lemma"* [HILL99] the simulated random oracle output is statistically close to uniform.
- Whenever $\mathsf{A}$ makes a signing query on $m$, $\mathsf{B}$ retrieves $(m, s_m)$ from the hash table and returns $s_m$ as the signature. Again it can be shown that the distribution of the signature is statistically close to the real one.
- Finally, upon receiving a forgery from $\mathsf{A}$, $\mathsf{B}$ can use this to find a collision. That is, two different preimages that map to the same hash value giving a solution to SIS. When $\mathsf{A}$ produces the forgery $(m^\star, s^\star)$, $\mathsf{B}$ looks up $(m^\star, s_{m^\star})$ in its hash table and outputs $(s^\star - s_{m^\star})$ as a SIS solution to $\boldsymbol{A}$. This is a valid solution to SIS because $\mathsf{H}(m^\star) = \boldsymbol{A}s^\star \mod q$ and $\mathsf{H}(m^\star) = \boldsymbol{A}s_{m^\star} \mod q$, and $\boldsymbol{A}(s^\star - s_{m^\star}) = 0 \mod q$ and $\|s^\star - s_{m^\star}\|$ is small.

*Plain unforgeability* can also be proven with a reduction to one-wayness (inhomogenous SIS). This proof is looser but enjoys better SIS parameters.

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 $\boldsymbol{h} \in \mathcal{R} = \mathbb{Z}_q[X]/(X^n + 1)$. Falcon instantiates a randomized version of the GPV framework with the NTRU-based preimage sampleable trapdoor $f_{\boldsymbol{h}}$ [HPS98, DLP14, PFH+22]. Specifically, $f_{\boldsymbol{h}}$ maps two ring elements $(\boldsymbol{s}_1, \boldsymbol{s}_2)$ to $\boldsymbol{s}_1 + \boldsymbol{h} * \boldsymbol{s}_2 \mod q$. Observe that $f_{\boldsymbol{h}}$ is a special case of the GPV trapdoor function $f_{\boldsymbol{A}}(s) = \boldsymbol{A}s \bmod q$. A valid signature on message $m$ consists of a tuple $(\boldsymbol{s}_1, \boldsymbol{s}_2) \in \mathcal{R}^2$ and a random salt $r \in \{0,1\}^k$ satisfying

$$\mathsf{H}(m, r) = \boldsymbol{s}_1 + \boldsymbol{h} * \boldsymbol{s}_2 \mod q \quad \wedge \quad \|(\boldsymbol{s}_1, \boldsymbol{s}_2)\|_2 \leq \beta.$$

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 in Falcon compared to the GPV framework is that signatures are not directly output from the preimage sampling procedure, as they may fail verification if their norms are too large – something that occurs with small probability of about $2^{-14}$. To eliminate this correctness error, signatures are checked for shortness, and if the norm exceeds some threshold, a new preimage is sampled repeatedly until one with a sufficiently small norm is found. This introduces a complication for simulating signing queries, as the process involves conditional distributions. The signing oracle outputs preimages conditioned on having sufficiently small norm, whereas programming the random oracle with this constraint and arguing about the uniformity of outputs seems 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 includes these changes [Por25a, Por25b].

Rényi Divergence in Falcon. Another issue is that Falcon relies on the Rényi divergence, whereas the GPV framework uses the statistical distance to prove the closeness of the sampler and a Gaussian. Citing [Pre17, Lem. 6] and 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 the 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 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 $\mathsf{H}(m, r) := \boldsymbol{s}_1 + \boldsymbol{h} * \boldsymbol{s}_2 \mod q$ is Rényi close to uniform for Gaussian $\boldsymbol{s}_1, \boldsymbol{s}_2$. 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, $Q_{\mathsf{H}} = 2^{96} \gg Q_s$. 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. However, for a reduction to collision resistance we also need to program further queries which have to include the query used in the forgery output. We use a recent technique from [BRTZ24] to program a predefined amount of queries such that we can choose the optimal trade-off between the Rényi loss and the guessing loss.

Concrete Security of Falcon$^+$. Table 2 summarises the concrete security bounds of Falcon$^+$, our modified version of Falcon. Here $\mathcal{R}$-**SIS** denotes the ring version of **SIS** and $t$-$\mathcal{R}$-**ISIS** denotes the multi-target version of $\mathcal{R}$-**SIS** with $t$ targets. A key observation is that achieving *strong unforgeability* requires a norm bound of $2\beta$ when reducing to **SIS**, while *plain unforgeability* requires only a norm bound of $\beta$ in the reduction to **ISIS**. Since both **SIS** and **ISIS** become easier with larger values of $\beta$, this explains why neither Falcon$^+$-512 nor Falcon$^+$-1024 satisfies *strong unforgeability*, as indicated in Table 1.

The parameters of Falcon have been carefully chosen so that the Rényi bound $r_u^{Q_s} = (1 + \delta_u)^{Q_s}$ is a small constant for $Q_s = 2^{64}$ signing queries. However, the Rényi bound $r_u^{Q_{\mathsf{H}}} = (1 + \delta_u)^{Q_{\mathsf{H}}}$ would become extremely large for $Q_{\mathsf{H}} = 2^{96} \gg 2^{64}$ random oracle queries. When the random oracle is programmed for each of the $Q_{\mathsf{H}}$ queries (setting $\mathsf{H}(m, r) = \boldsymbol{s}_1 + \boldsymbol{h} * \boldsymbol{s}_2$ for known $(\boldsymbol{s}_1, \boldsymbol{s}_2)$), this leads to a loose Rényi bound. This allows the reduction to construct a collision from one of the programmed preimages and the forgery. Additionally, the RO must be programmed for each signing query, but this can be neglected since $Q_{\mathsf{H}} \gg Q_s$.

Theorem 1 improves on this approach by leveraging techniques from [BRTZ24], where not every random oracle query is programmed. Instead, only a predefined number $L \in [Q_{\mathsf{H}}]$ queries are programmed. This reduces the loss to $r_u^{Q_s+L}$. Since the reduction involves programming the correct query, there is an additional guessing loss of $Q_{\mathsf{H}}/L$. Theorem 2 proves only plain unforgeability but reducing to multi-target $\mathcal{R}$-**ISIS** ($Q_{\mathsf{H}}$-$\mathcal{R}$-**ISIS**) instead of $\mathcal{R}$-**SIS**. In this reduction, a smaller norm bound is required. The challenge targets are embedded directly into the random oracle, and one of them can be solved upon receiving a forgery. As a result, the distribution of the random oracle must only be changed for the signing queries resulting in a Rényi loss of $r_u^{Q_s}$.

The resulting bit security for Falcon$^+$-512 (NIST Level I) and Falcon$^+$-1024 (NIST Level V) are shown in Table 1 on Page 4. 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 $\mathcal{R}$-**(I)SIS**.

| **Security** | **Multiplicative loss** | **Assumption** |
|---|---|---|
| **SUF-CMA** (Th. 1) | $\frac{Q_{\mathsf{H}}}{L} \cdot r_u^{Q_s+L} \cdot r_p^{Q_s}$ | $\mathcal{R}$-**SIS**$_{2\beta}$ |
| **UF-CMA** (Th. 2) | $r_u^{Q_s} \cdot r_p^{Q_s}$ | $Q_{\mathsf{H}}$-$\mathcal{R}$-**ISIS**$_\beta$ |

**Table 2.** Concrete security loss (simplified) for Falcon$^+$ in the random oracle model. Constants $r_u = 1 + \delta_u$ and $r_p = 1 + \delta_p$ are Rényi divergences related to the uniformity of an NTRU evaluation on Gaussian inputs ($r_u$) and the preimage sampler ($r_p$). $Q_s$ and $Q_{\mathsf{H}}$ denote the number of signing and random oracle queries, respectively. $L \in [Q_{\mathsf{H}}]$ denotes the number of programmed random oracle queries and describes a trade-off between guessing loss and Rényi loss.

### 1.3 Open Problems

We do not present specific attacks against Falcon, nor do we claim that a stronger security proof is impossible. It may be possible to improve the analysis, and we consider this an open problem worth exploring. Falcon relies on [Pre17, Lem. 6] and the analysis of the Klein sampler [Kle00], to claim that the Rényi divergence between the FFO sampler's output and an ideal Gaussian is small. An important area for future work is to conduct a similar analysis for the FFO sampler as well as the key generation procedure. Finally, we leave open as an interesting future direction the potential application of the Rényi divergence in the quantum random oracle model (QROM). We remark that Hawk [BBD$^+$23] was analysed in the QROM [FH23] but its proof only requires programming $Q_s$ fixed positions in the QROM and no direct queries, unlike ours. Furthermore, they do not need to program the random oracle beyond responding to signing queries because they reduce to a different, "one-more"-type hardness assumption.

## 2 Preliminaries

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

### 2.1 Notations

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 $[n] := \{1, \ldots, n\}$. The notation $[\![b]\!]$, where $b$ is a boolean statement, evaluates to 1 if the statement is true and 0 otherwise. We use uppercase letters $\mathsf{A}, \mathsf{B}, \mathsf{C}, \mathsf{D}$ to denote algorithms. Unless otherwise stated, algorithms are probabilistic, and we write $(y_1, \ldots) \xleftarrow{\$} \mathsf{A}(x_1, \ldots)$ to denote that $\mathsf{A}$ returns $(y_1, \ldots)$ when run on input $(x_1, \ldots)$. We write $\mathsf{A}^{\mathsf{B}}$ to denote that $\mathsf{A}$ has oracle access to $\mathsf{B}$ during its execution. The support of a discrete random variable $X$ is defined as $\sup(X) := \{x \in \mathbb{R} \mid \Pr[X = x] > 0\}$. For two polynomials $\boldsymbol{f}, \boldsymbol{g}$, we denote the polynomial multiplication of $\boldsymbol{f}$ and $\boldsymbol{g}$ by $\boldsymbol{f} * \boldsymbol{g}$. 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 [BR04]. A *game* $\mathsf{G}$ is a probability experiment in which an adversary $\mathsf{A}$ interacts with an implicit challenger that answers oracle queries issued by $\mathsf{A}$. The game $\mathsf{G}$ 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 $\mathsf{G}$ between a challenger and an adversary $\mathsf{A}$ as $\mathsf{G}(\mathsf{A}) \Rightarrow b$. $\mathsf{A}$ is said to *win* $\mathsf{G}$ if $\mathsf{G}^{\mathsf{A}} \Rightarrow 1$, or shortly $\mathsf{G} \Rightarrow 1$. Unless otherwise stated, the randomness in the probability term $\Pr[\mathsf{G}(\mathsf{A}) \Rightarrow 1]$ is over all the random coins in game $\mathsf{G}$. To provide a cleaner description and avoid repetitions, we sometimes refer to procedures of different games. To call the oracle procedure $\texttt{Oracle}$ of game $\mathsf{G}$ on input $x$, we shortly write $\mathsf{G}.\texttt{Oracle}(x)$. 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.

### 2.2 Signatures

We recall the syntax and standard security notions of signatures.

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

$(sk, pk) \xleftarrow{\$} \mathsf{Gen}$**:** The probabilistic key generation algorithm returns a secret key $sk$ and a corresponding public key $pk$, where $pk$ defines a message space $\mathcal{M}$.

$\sigma \xleftarrow{\$} \mathsf{Sgn}(sk, m)$**:** Given a secret key $sk$ and a message $m \in \mathcal{M}$, the probabilistic signing algorithm $\mathsf{Sgn}$ returns a signature $\sigma$.

$b \leftarrow \mathsf{Ver}(pk, m, \sigma)$**:** Given a public key $pk$, a message $m$, and a signature $\sigma$, the deterministic verification algorithm $\mathsf{Ver}$ returns a bit $b$, such that $b = 1$ if and only if $\sigma$ is a valid signature on $m$ and $b = 0$ otherwise.

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

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

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

### 2.3 Lattices

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

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

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

```
Games $Q_s$-UF-CMA$_{\mathsf{Sig}}(\mathsf{A})/Q_s$-SUF-CMA$_{\mathsf{Sig}}(\mathsf{A})$
$\mathcal{Q} \leftarrow \emptyset$
$(sk, pk) \xleftarrow{\$} \mathsf{Gen}$
$(m^\star, \sigma^\star) \xleftarrow{\$} \mathsf{A}^{\mathtt{Sgn}(\cdot)}(pk)$
return $[\![\mathsf{Ver}(m^\star, \sigma^\star) = 1 \land (m^\star, \cdot) \notin \mathcal{Q}]\!]$        // UF-CMA
return $[\![\mathsf{Ver}(pk, m^\star, \sigma^\star) = 1 \land (m^\star, \sigma^\star) \notin \mathcal{Q}]\!]$    // SUF-CMA

Oracle $\mathtt{Sgn}(m)$
$\sigma \xleftarrow{\$} \mathsf{Sgn}(sk, m)$
$\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(m, \sigma)\}$
return $\sigma$
```

**Figure 2.** Games defining **UF-CMA** and **SUF-CMA** for a signature scheme $\mathsf{Sig} = (\mathsf{Gen}, \mathsf{Sgn}, \mathsf{Ver})$ and adversary $\mathsf{A}$ making at most $Q_s$ queries to $\mathtt{Sgn}$.

Anticirculant matrices satisfy the following useful properties.

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

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 $\boldsymbol{f}$ instead of $\mathcal{A}(\boldsymbol{f})$.

Let the $\ell_2$-norm for $\boldsymbol{f} = f_0 + f_1 X + \ldots + f_{n-1} X^{n-1} \in \mathcal{R}_q$ be defined as $\|\boldsymbol{f}\|_2 \coloneqq \sqrt{\sum_{i=0}^{n-1} |f_i|^2}$. For two polynomials $\boldsymbol{f}, \boldsymbol{g} \in \mathcal{R}_q$ we use the notation

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

Lattices. A lattice $\boldsymbol{\Lambda} \subseteq \mathbb{R}^n$ 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 $b_1, \ldots, b_m \in \mathbb{R}^n$ of *linearly independent* vectors that form a basis $\boldsymbol{B} = \{b_1, \ldots, b_m\}$ for the lattice

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

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

The *determinant* of a lattice $\boldsymbol{\Lambda} = \boldsymbol{\Lambda}(\boldsymbol{B}) \subseteq \mathbb{R}^n$ for some basis $\boldsymbol{B} \in \mathbb{R}^{n \times m}$ is defined as $\det(\boldsymbol{\Lambda}) = \sqrt{\det(\boldsymbol{B}^\top \boldsymbol{B})}$. The *orthogonal* lattice for $\boldsymbol{A} \in \mathbb{Z}_q^{n \times m}$ is defined as $\boldsymbol{\Lambda}^\perp(\boldsymbol{A}) \coloneqq \{e \in \mathbb{Z}^m \mid \boldsymbol{A}e = 0 \mod q\}$. For an $n$-dimensional lattice $\boldsymbol{\Lambda}$, a lattice $\boldsymbol{\Lambda}' \subseteq \boldsymbol{\Lambda}$ is called a sublattice of $\boldsymbol{\Lambda}$. One can define the following quotient group $\boldsymbol{\Lambda}/\boldsymbol{\Lambda}' \coloneqq \{v + \boldsymbol{\Lambda}' \mid v \in \boldsymbol{\Lambda}\}$, which forms a group under the addition of cosets $v + \boldsymbol{\Lambda}'$.

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

$$\boldsymbol{\Lambda}_{\boldsymbol{h},q} \coloneqq \{(\boldsymbol{u}, \boldsymbol{v}) \in \mathcal{R}^2 \mid \boldsymbol{u} + \boldsymbol{v} * \boldsymbol{h} = \boldsymbol{0} \mod q\}.$$

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

$$\boldsymbol{B}_{\boldsymbol{h}} = \begin{bmatrix} -\boldsymbol{h} & \boldsymbol{q} \\ \boldsymbol{1} & \boldsymbol{0} \end{bmatrix}$$

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

$$\boldsymbol{B}_{\boldsymbol{f},\boldsymbol{g}} = \begin{bmatrix} \boldsymbol{g} & \boldsymbol{G} \\ -\boldsymbol{f} & -\boldsymbol{F} \end{bmatrix}$$

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

Gaussians and Preimage Sampling. We define discrete Gaussians and state some of their useful properties.

**Definition 6 (Discrete Gaussian Distribution over $\boldsymbol{\Lambda}$).** The $n$-dimensional *Gaussian function* $\rho_{s,c} \colon \mathbb{R}^n \to (0, 1]$ on $\mathbb{R}^n$ centered at $c \in \mathbb{R}^n$ with standard deviation $s > 0$ is defined by

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

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

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

We sometimes use the following notation $\rho_{s,c}(\boldsymbol{\Lambda}) = \sum_{x \in \boldsymbol{\Lambda}} \rho_{s,c}(x)$. We omit the subscript $c$ when the Gaussian is centered at 0 and subscript $\boldsymbol{\Lambda}$ when the Gaussian is over $\mathbb{Z}^n$. We use $\boldsymbol{f} \sim \mathcal{D}_{\mathcal{R}}$ to denote the polynomial $\boldsymbol{f} := \sum_{i=0}^{n-1} f_i X^i \mod (X^n + 1)$ 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 [DRSD14, Lem. 2.13] [ADRS15, Lem. 2]).** For any lattice $\boldsymbol{\Lambda} \in \mathbb{R}^n$, $s > 0$, $c \in \mathbb{R}^n$, and $\tau > 1$,

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

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

**Lemma 3 (NTRU Trapdoor Generation [HPS98, Pre15]).** An NTRU Trapdoor Generation algorithm $\mathsf{TpdGen}(\mathcal{R}, \alpha, q)$, given a ring $\mathcal{R}$, a target quality $\alpha \geq 1$, and a modulus $q$, returns a public key $\boldsymbol{h} \in \mathcal{R}_q \setminus \{\mathbf{0}\}$ and the trapdoor $(\boldsymbol{f}, \boldsymbol{g}, \boldsymbol{F}, \boldsymbol{G}) \in \mathcal{R}^4$, such that $\boldsymbol{B}_{\boldsymbol{h}}$ and $\boldsymbol{B}_{\boldsymbol{f},\boldsymbol{g}}$ form a basis of the same lattice. Furthermore, $\|\boldsymbol{B}_{\boldsymbol{f},\boldsymbol{g}}\|_{GS} \leq \alpha\sqrt{q}$.

Let $\boldsymbol{\Lambda}$ be an $n$-dimensional lattice and $\epsilon > 0$, the (scaled) smoothing parameter $\eta_\epsilon(\boldsymbol{\Lambda})$ is the smallest $s > 0$ such that $\rho_{1/s}(\boldsymbol{\Lambda}^* \setminus 0) \leq \epsilon$, where $\boldsymbol{\Lambda}^*$ 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. 4.4]).** For any $\epsilon \in (0, 1)$ it holds that

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

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

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

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

**Lemma 6 (Min-Entropy of Gaussian (implicit in [PR06, Lem. 2.10])).** Let $n \in \mathbb{N}$, $\mathbf{\Lambda} = \mathbb{Z}^n$, $\epsilon \in (0, \frac{1}{2})$ and $s \geq \eta_\epsilon(\mathbf{\Lambda}) > 2$. Then for any $c \in \mathbb{R}^n$ and $x \in \mathbf{\Lambda}$,

$$\mathcal{D}_{\mathbf{\Lambda},s,c}(x) \leq \frac{1}{2^{n-1}}.$$

The proof can be found in Appendix A.1

### 2.4 Rényi Divergence

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

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

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

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

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 $\delta$ indicates that the distributions are closer.

**Lemma 7 (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, $R_a(\mathcal{P} \mid\mid \mathcal{Q}) = R_a(\mathcal{P}_1 \mid\mid \mathcal{Q}_1) \cdot R_a(\mathcal{P}_2 \mid\mid \mathcal{Q}_2)$.

**Lemma 8 (Probability Preservation [LSS14, Lem. 4.1]).** Let $a \in (1, \infty)$ and $E \subseteq \sup(\mathcal{Q})$ be an arbitrary event. Then,

$$\mathcal{Q}(E) \geq \mathcal{P}(E)^{\frac{a}{a-1}} / R_a(\mathcal{P} \mid\mid \mathcal{Q})$$
$$\mathcal{Q}(E) \geq \mathcal{P}(E) / R_\infty(\mathcal{P} \mid\mid \mathcal{Q}).$$

**Lemma 9 (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 $x \leftarrow \mathcal{Q}$), $R_a(\mathcal{P}^f \mid\mid \mathcal{Q}^f) \leq R_a(\mathcal{P} \mid\mid \mathcal{Q})$.

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

**Lemma 10 (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 $(X_i)_{i<N}$. 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 varibles $(X_0, \ldots, X_{i-1}) = X_{<i}$. 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,

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

Then,

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

**Definition 9 (Relative Error [MW17]).** *Let $\mathcal{P}$ and $\mathcal{Q}$ be two discrete probability distributions over the same countable set $\mathcal{X}$. The relative error of $\mathcal{P}$ and $\mathcal{Q}$ is defined as*

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

The relative error can be used to bound the Renyi Divergence.

**Lemma 11 (Relative Error [Pre17, Lem. 3]).** Let $\mathcal{P}, \mathcal{Q}$ be two distributions such that $\sup(\mathcal{P}) = \sup(\mathcal{Q})$ and $\delta_{RE} > 0$. Then for all $a \in (1, +\infty)$

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

The Klein Sampler [Kle00, GPV08] was analyzed in [Pre17] with respect to its relative error and Rényi divergence.

**Lemma 12 (Relative Error of Klein Sampler [Pre15, Pre17]).** Let $n$ be a positive integer and $\epsilon \in (0, 1/4)$. Then the *relative error* of the Klein Sampler $\mathsf{PreSmp}(\boldsymbol{B}, s, (\boldsymbol{c}, \boldsymbol{0}))$ for any basis $\boldsymbol{B} \in \mathbb{Z}^{2n \times 2n}$, standard deviation $s \geq \eta_\epsilon(\mathbb{Z}^{2n}) \cdot \|\boldsymbol{B}\|_{GS}$, and arbitrary syndrome $\boldsymbol{c} \in \mathcal{R}_q$ is bounded by

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

**Corollary 1 (Rényi Divergence of Klein Sampler).** Let $n$ be a be a positive integer, $a > 1$, and $\epsilon \in (0, 1/4)$. Then for the Klein Sampler $\mathsf{PreSmp}(\boldsymbol{B}, s, (\boldsymbol{c}, \boldsymbol{0}))$, for any basis $\boldsymbol{B} \in \mathbb{Z}^{2n \times 2n}$, standard deviation $s \geq \eta_\epsilon(2^{2n}) \cdot \|\boldsymbol{B}\|_{GS}$, and arbitrary syndrome $\boldsymbol{c} \in \mathcal{R}_q$, the *Rényi divergence* is bounded by

$$R_a\left(\mathsf{PreSmp}(\boldsymbol{B}, s, (\boldsymbol{c}, \boldsymbol{0})) \mid\mid \mathcal{D}_{\boldsymbol{\Lambda}(\boldsymbol{B}), s, (\boldsymbol{c}, \boldsymbol{0})}\right) \lesssim 1 + 2a\epsilon^2.$$

### 2.5 Hardness Assumptions

We will use the following definitions of the $\mathcal{R}$-**SIS** and $t$-$\mathcal{R}$-**ISIS** problems over NTRU lattices. In the multi-target version of $\mathcal{R}$-**ISIS**, the adversary is given $t$ $\mathcal{R}$-**ISIS** challenges and can choose one to solve.

**Definition 10 ($\mathcal{R}$-SIS, $t$-$\mathcal{R}$-ISIS).** The *Ring Short Integer Solution* problem and the *Ring Inhomogeneous Short Integer Solution* problem relative to the NTRU trapdoor algorithm $\mathsf{TpdGen}$ with parameters $m, q > 0$ and $\alpha, B > 0$ are defined via the games $\mathcal{R}$-**SIS** and $t$-$\mathcal{R}$-**ISIS**, depicted in Figure 3. We define the advantages of $\mathsf{A}$ as

$$\begin{aligned} \mathrm{Adv}^{\mathcal{R}\text{-}\mathbf{SIS}}_{m,q,\alpha,B,\mathsf{A}} &\coloneqq \Pr[\mathcal{R}\text{-}\mathbf{SIS}_{m,q,\alpha,B}(\mathsf{A}) \Rightarrow 1], \\ \mathrm{Adv}^{t\text{-}\mathcal{R}\text{-}\mathbf{ISIS}}_{m,q,\alpha,B,\mathsf{A}} &\coloneqq \Pr[t\text{-}\mathcal{R}\text{-}\mathbf{ISIS}_{m,q,\alpha,B}(\mathsf{A}) \Rightarrow 1]. \end{aligned}$$

According to [LM06], $\mathcal{R}$-$\mathbf{SIS}_{m,q,\alpha,B}$ and 1-$\mathcal{R}$-$\mathbf{ISIS}_{m,q,\alpha,B}$ are as hard as $\mathbf{SVP}_\gamma$ for $\gamma = \tilde{O}(nB)$. In particular, its hardness is independent of $m$. Note that we defined **(I)SIS** with respect to an NTRU key instead of a uniformly random element since **(I)SIS** is not believed to become easier in this 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 **(I)SIS** definition. We make the assumption that $\mathcal{R}$-**SIS** and $t$-$\mathcal{R}$-**ISIS** instances are as hard as random **SIS** and **ISIS** instances. Although we note that there may exist better attacks against $t$-$\mathcal{R}$-**ISIS** [Ber22].

```
Game $\mathcal{R}$-$\mathbf{SIS}_{m,q,\alpha,B}(\mathsf{A})$
for $i \in [m]$ :
    $(\boldsymbol{h}_i, \cdot, \cdot) \leftarrow \mathsf{TpdGen}(q, \alpha)$
$(\boldsymbol{u}, \boldsymbol{v}_1, \ldots, \boldsymbol{v}_m) \xleftarrow{\$} \mathsf{A}(\boldsymbol{h}_1, \ldots, \boldsymbol{h}_m)$
return $[\![\boldsymbol{u} + \sum_{i \in [m]} \boldsymbol{h}_i * \boldsymbol{v}_i = 0 \wedge 0 < \|(\boldsymbol{u}, \boldsymbol{v}_1, \ldots, \boldsymbol{v}_m)\|_2 \leq B]\!]$
Game $t$-$\mathcal{R}$-$\mathbf{ISIS}_{m,q,\alpha,B}(\mathsf{A})$
for $i \in [m]$ :
    $(\boldsymbol{h}_i, \cdot, \cdot) \leftarrow \mathsf{TpdGen}(q, \alpha)$
for $j \in [t]$ :
    $\boldsymbol{c}_j \xleftarrow{\$} \mathcal{R}_q$
$(j, \boldsymbol{u}, \boldsymbol{v}_1, \ldots, \boldsymbol{v}_m) \xleftarrow{\$} \mathsf{A}(\boldsymbol{h}_1, \ldots, \boldsymbol{h}_m, \boldsymbol{c}_1, \ldots, \boldsymbol{c}_t)$
return $[\![\boldsymbol{u} + \sum_{i \in [m]} \boldsymbol{h}_i * \boldsymbol{v}_i = \boldsymbol{c}_j \wedge \|(\boldsymbol{u}, \boldsymbol{v}_1, \cdots, \boldsymbol{v}_m)\|_2 \leq B]\!]$
```

**Figure 3.** Games defining $\mathcal{R}$-$\mathbf{SIS}_{m,q,\alpha,B}$ and $t$-$\mathcal{R}$-$\mathbf{ISIS}_{m,q,\alpha,B}$.

## 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 $\infty$ 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 13 shows that a Gaussian sample over $\boldsymbol{\Lambda}$ is distributed almost-uniformly modulo a sublattice $\boldsymbol{\Lambda}'$, provided the standard deviation exceeds the smoothing parameter of $\boldsymbol{\Lambda}'$.

**Lemma 13 (Rényi Divergence of Gaussian Sample over $\boldsymbol{\Lambda}/\boldsymbol{\Lambda}'$ (adapted from [GPV08, Cor. 2.8])).** Let $\boldsymbol{\Lambda}, \boldsymbol{\Lambda}'$ be $n$-dimensional full-rank lattices with $\boldsymbol{\Lambda}' \subseteq \boldsymbol{\Lambda}$. Then for any $a \in (1, \infty)$, $\epsilon \in (0, \frac{1}{2})$, any $s \geq \eta_\epsilon(\boldsymbol{\Lambda}')$, and any $c \in \mathbb{R}^n$,

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

The proof can be found in Appendix A.2. 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 14 (Rényi divergence (adapted from [GPV08, Lem 5.2])).** If the columns of $\boldsymbol{A} \in \mathbb{Z}_q^{n \times m}$ generate $\mathbb{Z}_q^n$, $a \in (1, \infty)$, $\epsilon \in (0, \frac{1}{2})$, and $s \geq \eta_\epsilon(\boldsymbol{\Lambda}^\perp(\boldsymbol{A}))$; 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 $u = \boldsymbol{A}e \mod q$, it holds that

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

The proof can be found in Appendix A.3.

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

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

The proof can be found in Appendix A.4.

The next lemma shows that the tailbounds of two distributions with a small relative error are close.[4]

**Lemma 15 (Relative Error for Tailbounds).** Let $\mathcal{P}$ and $\mathcal{Q}$ be two distributions with $\sup(\mathcal{P}) = \sup(\mathcal{Q}) = \mathbb{Z}^n$ such that their relative error is bounded by $\frac{\mathcal{P}}{\mathcal{Q}} \leq 1 + \delta$. 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 Appendix A.5.

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 16 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 16 (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 $Q \in \mathbb{N}$ and $a \in (1, \infty)$ the Rényi divergence between two arbitrary distributions is at most $R_a \in [1, \infty)$, and

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

Then

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

The proof can be found in Appendix A.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 $\mathsf{TpdGen} : \mathcal{R} \times \mathbb{R} \times \mathbb{Z} \to \mathcal{R}^4$ be a trapdoor generation algorithm, let $\mathsf{PreSmp} : \mathbb{Z}^{2n \times 2n} \times \mathbb{R} \times \mathcal{R}^2 \to \mathcal{R}^2$ be a preimage sampling algorithm, and $\mathsf{H} : \mathcal{R}_q \times \{0,1\}^k \times \mathcal{M} \to \mathcal{R}_q$ be a hash function. The defining algorithms of signature schemes CoreFalcon$^+$ and CoreFalcon are given in Figure 4.

Note that CoreFalcon$^+$ is a slight modification of CoreFalcon: In signing $\mathsf{Sgn}^+$ of CoreFalcon$^+$, picking the random seed $r$ and computing the ring element $\boldsymbol{c} = \mathsf{H}(pk, r, m)$ is performed inside the repeat loop (lines 15-19), 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.[5] Unfortunately, we are 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].

[4] Note that the relative error is related to Rényi arguments via Lemma 11.

[5] 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].

```
Gen
$(\boldsymbol{f},\boldsymbol{g},\boldsymbol{F},\boldsymbol{G}) \xleftarrow{\$} \mathsf{TpdGen}(\mathcal{R},\alpha,q)$
$\boldsymbol{B} := \left[\begin{array}{c|c} \mathcal{A}(\boldsymbol{g}) & \mathcal{A}(\boldsymbol{G}) \\ \hline -\mathcal{A}(\boldsymbol{f}) & -\mathcal{A}(\boldsymbol{F}) \end{array}\right] \in \mathbb{Z}^{2n\times 2n}$
$\boldsymbol{h} := \boldsymbol{g} * \boldsymbol{f}^{-1} \in \mathcal{R}_q$
return $(sk := \boldsymbol{B}, pk := \boldsymbol{h})$

Sgn$(sk = \boldsymbol{B}, m)$
$r \xleftarrow{\$} \{0,1\}^k$
$\boldsymbol{c} := \mathsf{H}(pk, r, m) \in \mathcal{R}_q$
repeat
    $(\boldsymbol{s}_1, \boldsymbol{s}_2) \xleftarrow{\$} \mathsf{PreSmp}(\boldsymbol{B}, s, (\boldsymbol{c}, \boldsymbol{0}))$
until $\|(\boldsymbol{s}_1, \boldsymbol{s}_2)\|_2 \leq \beta$
$\sigma := (r, \boldsymbol{s}_2) \in \{0,1\}^k \times \mathcal{R}$
return $\sigma$

Ver$(pk = \boldsymbol{h}, m, \sigma = (r, \boldsymbol{s}_2))$
$\boldsymbol{c} := \mathsf{H}(pk, r, m)$
$\boldsymbol{s}_1 := \boldsymbol{c} - \boldsymbol{s}_2 * \boldsymbol{h} \mod q$
return $[\![\|(\boldsymbol{s}_1, \boldsymbol{s}_2)\|_2 \leq \beta]\!]$

Sgn$^+(sk = \boldsymbol{B}, m)$
repeat
    $r \xleftarrow{\$} \{0,1\}^k$
    $\boldsymbol{c} := \mathsf{H}(pk, r, m) \in \mathcal{R}_q$
    $(\boldsymbol{s}_1, \boldsymbol{s}_2) \xleftarrow{\$} \mathsf{PreSmp}(\boldsymbol{B}, s, (\boldsymbol{c}, \boldsymbol{0}))$
until $\|(\boldsymbol{s}_1, \boldsymbol{s}_2)\|_2 \leq \beta$
$\sigma := (r, \boldsymbol{s}_2) \in \{0,1\}^k \times \mathcal{R}$
return $\sigma$
```

**Figure 4.** Construction of the CoreFalcon = (Gen, Sgn, Ver) and CoreFalcon$^+$ = (Gen, Sgn$^+$, Ver) signature schemes.

### 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 $\lambda$ is the security parameter, set to 128 for NIST level I and 256 for NIST level V. Given $\epsilon$, the standard deviation $s$ is given by

$$s = \frac{1}{\pi}\sqrt{\frac{\ln(4n(1+1/\epsilon))}{2}} \cdot 1.17\sqrt{q}.$$

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 $\beta$ 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 3. We define Falcon$^+$-512 and Falcon$^+$-1024 using the CoreFalcon$^+$ framework, instantiated with the parameters from Table 3.

Falcon uses the FFO sampler to instantiate the preimage sampler PreSmp. Since a formal analysis of the FFO sampler is lacking, we base our analysis and security estimation on the Klein sampler's [Kle00] analysis from [Pre17] which is expected to closely approximate the FFO sampler. We stress that future work is needed to analyse the FFO sampler, which may slightly alter the results presented here. However, this only affects the parameter selection, in particular the optimisation of the Rényi order. The rest of our analysis remains unchanged, as it is modular, general, and parameterised by the sampler.

### 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 *strong* unforgeability. This result leverages a proof technique introduced in [BRTZ24], which defines a parameter $L$ that enables a trade-off between the tightness loss from a guessing argument and the Rényi loss. Theorem 2 provides only plain unforgeability, but is based on a weaker assumption.

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

**Table 3.** Parameter sets for Falcon-512/Falcon$^+$-512 and Falcon-1024/Falcon$^+$-1024 [PFH$^+$22, Tab. 3.3].

**Theorem 1 (Strong Unforgeability).** For any adversary $\mathsf{A}$ against the **SUF-CMA** security of CoreFalcon$^+$ (Figure 4) running in time $t_{\mathsf{A}}$, making at most $Q_s$ signing queries and $Q_{\mathsf{H}}$ random oracle queries, there exist an adversary $\mathsf{B}$ against $\mathcal{R}$-**SIS** running in time $t_{\mathsf{B}} \approx t_{\mathsf{A}}$ such that for all $C_s \in \mathbb{N}^{\geq 1}, L \in [Q_{\mathsf{H}}+1]$ and $a_u, a_p \in \mathbb{R}^{>1}$ it holds

$$\mathrm{Adv}^{Q_s\text{-}\mathbf{SUF\text{-}CMA}}_{\textsc{CoreFalcon}^+,\mathsf{A}} \leq \frac{(Q_{\mathsf{H}}+1)}{L \cdot p_{\mathsf{PreSmp},\beta}} \cdot \left( r_u^{C_s+L} \cdot \left( r_p^{C_s+1} \cdot \left( \mathrm{Adv}^{\mathcal{R}\text{-}\mathbf{SIS}}_{1,q,\alpha,2\beta,\mathsf{B}} + \frac{1}{2^{n-1}} \right) \right)^{\frac{a_p-1}{a_p}} \right)^{\frac{a_u-1}{a_u}} + \sum_{i=0}^{Q_s} \binom{C_s}{i} (1-p_{\mathsf{PreSmp},\beta})^{C_s-i} (p_{\mathsf{PreSmp},\beta})^i + \frac{Q_s(C_s+Q_{\mathsf{H}})}{2^k} \ ,$$

where

$p_{\mathsf{PreSmp},\beta} := \max_{(\boldsymbol{B},\cdot)\in\mathsf{TpdGen},\boldsymbol{c}\in\mathcal{R}_q} \Pr_{(\boldsymbol{s}_1,\boldsymbol{s}_2)\xleftarrow{\$}\mathsf{PreSmp}(\boldsymbol{B},s,(\boldsymbol{c},\boldsymbol{0}))}[\|(\boldsymbol{s}_1,\boldsymbol{s}_2)\|_2 \leq \beta]$,
$r_u = \max_{(\cdot,\boldsymbol{h})\in\mathsf{TpdGen}} R_{a_u}(\mathcal{P} \,||\, \mathcal{Q}_{\boldsymbol{h}})$ with $\mathcal{P} = \mathcal{U}(\mathcal{R}_q)$ and $\mathcal{Q}_{\boldsymbol{h}}$ the distribution of $\boldsymbol{s}_1 + \boldsymbol{s}_2 * \boldsymbol{h} \mod q$, where $\boldsymbol{s}_1, \boldsymbol{s}_2 \sim \mathcal{D}_{\mathcal{R},s}$,
$r_p = \max_{(\boldsymbol{B},\cdot)\in\mathsf{TpdGen},\boldsymbol{c}\in\mathcal{R}_q} R_{a_p}(\mathsf{PreSmp}(\boldsymbol{B},s,(\boldsymbol{c},\boldsymbol{0})) \,||\, \mathcal{D}_{\Lambda(\boldsymbol{B},s,(\boldsymbol{c},\boldsymbol{0}))})$.

The proof of Theorem 1 can be found in Appendix B.

Remark. Note that the bound of Theorem 1 holds for all choices of constants $C_s \in \mathbb{N}^{\geq 1}$, $L \in [Q_{\mathsf{H}}+1]$, and $a_u, a_p \in \mathbb{R}^{>1}$. We will refer to them as *proof constants*. Later, in Section 6 we will derive optimal choices for the proof constants that minimise the security loss for concrete and relevant instantiations of CoreFalcon$^+$.

**Theorem 2 (Unforgeability).** For any adversary $\mathsf{A}$ against the **UF-CMA** security of CoreFalcon$^+$ (Figure 4) running in time $t_{\mathsf{A}}$, making at most $Q_s$ signing queries and $Q_{\mathsf{H}}$ random oracle queries, there exist adversary $\mathsf{B}$ against $(Q_{\mathsf{H}}+1)$-$\mathcal{R}$-**ISIS** running in time $t_{\mathsf{B}} \approx t_{\mathsf{A}}$ such that for all $C_s \in \mathbb{N}^{\geq 1}$ and $a_u, a_p \in \mathbb{R}^{>1}$ it holds

$$\mathrm{Adv}^{Q_s\text{-}\mathbf{UF\text{-}CMA}}_{\textsc{CoreFalcon}^+,\mathsf{A}} \leq \left( r_u^{C_s} \cdot \left( r_p^{C_s} \cdot \mathrm{Adv}^{(\mathsf{Q_H}+1)\text{-}\mathcal{R}\text{-}\mathbf{ISIS}}_{1,q,\alpha,\beta,\mathsf{B}} \right)^{\frac{a_p-1}{a_p}} \right)^{\frac{a_u-1}{a_u}} + \sum_{i=0}^{Q_s} \binom{C_s}{i} (1-p_{\mathsf{PreSmp},\beta})^{C_s-i} (p_{\mathsf{PreSmp},\beta})^i + \frac{Q_s(C_s+Q_{\mathsf{H}})}{2^k} \ ,$$

where

$p_{\mathsf{PreSmp},\beta} := \max_{(\boldsymbol{B},\cdot)\in\mathsf{TpdGen},\boldsymbol{c}\in\mathcal{R}_q} \Pr_{(\boldsymbol{s}_1,\boldsymbol{s}_2)\xleftarrow{\$}\mathsf{PreSmp}(\boldsymbol{B},s,(\boldsymbol{c},\boldsymbol{0}))}[\|(\boldsymbol{s}_1,\boldsymbol{s}_2)\|_2 \leq \beta]$,

$r_u = \max_{(\cdot,\boldsymbol{h})\in\mathsf{TpdGen}} R_{a_u}(\mathcal{P} \mid\mid \mathcal{Q}_{\boldsymbol{h}})$ with $\mathcal{P} = \mathcal{U}(\mathcal{R}_q)$ and $\mathcal{Q}_{\boldsymbol{h}}$ the distribution of $\boldsymbol{s}_1 + \boldsymbol{s}_2 * \boldsymbol{h} \mod q$, where $\boldsymbol{s}_1, \boldsymbol{s}_2 \sim \mathcal{D}_{\mathcal{R},s}$,

$r_p = \max_{(\boldsymbol{B},\cdot)\in\mathsf{TpdGen},\boldsymbol{c}\in\mathcal{R}_q} R_{a_p}(\mathsf{PreSmp}(\boldsymbol{B},s,(\boldsymbol{c},\boldsymbol{0})) \mid\mid \mathcal{D}_{\Lambda(\boldsymbol{B},s,(\boldsymbol{c},\boldsymbol{0})})$.

The proof of Theorem 2 can be found in Section 5. Interestingly, Theorem 2 not only requires the hardness of $(Q_{\mathsf{H}}+1)$-$\mathcal{R}$-**ISIS**, but is, in fact, also necessary. Specifically, an attack on $(Q_{\mathsf{H}}+1)$-$\mathcal{R}$-**ISIS** would directly lead to an attack on Falcon.

## 5 Proof of Theorem 2

Consider the sequence of games depicted in Figure 5.

*Game* $\mathsf{G}_0$. This is the unforgeability game for CoreFalcon$^+$ so by definition we have

$$\Pr[\mathsf{G}_0^{\mathsf{A}} \Rightarrow 1] = \mathrm{Adv}^{Q_s\text{-}\mathbf{UF\text{-}CMA}}_{\textsc{CoreFalcon}^+,\mathsf{A}}.$$

```
Games $\mathsf{G}_0 - \mathsf{G}_5$
$\mathcal{H}, \mathcal{Q} \leftarrow \emptyset$
$cnt := 0$
$(\boldsymbol{B}, \boldsymbol{h}) \xleftarrow{\$} \mathsf{Gen}$
$(m^\star, \sigma^\star) \xleftarrow{\$} \mathsf{A}^{\mathtt{Sgn}(\cdot),\mathsf{H}(\cdot,\cdot,\cdot)}(\boldsymbol{h})$
return $[\![\mathsf{Ver}(\boldsymbol{h}, m^\star, \sigma^\star) = 1 \wedge (m^\star, \cdot) \notin \mathcal{Q}]\!]$

Oracle $\mathtt{Sgn}(m)$
repeat
    $cnt \leftarrow cnt + 1$                                        // $\mathsf{G}_1 - \mathsf{G}_5$
    if $cnt > C_s$                                                 // $\mathsf{G}_1 - \mathsf{G}_5$
        abort                                                      // $\mathsf{G}_1 - \mathsf{G}_5$
    $r \xleftarrow{\$} \{0,1\}^k$
    $\boldsymbol{c} := \mathsf{H}(\boldsymbol{h}, r, m)$                        // $\mathsf{G}_0 - \mathsf{G}_1$
    $(\boldsymbol{c}, \boldsymbol{s}_1, \boldsymbol{s}_2) := \mathsf{H}'(\boldsymbol{h}, r, m)$             // $\mathsf{G}_2 - \mathsf{G}_5$
    $(\boldsymbol{s}_1, \boldsymbol{s}_2) \xleftarrow{\$} \mathsf{PreSmp}(\boldsymbol{B}, s, (\boldsymbol{c}, \boldsymbol{0}))$      // $\mathsf{G}_0 - \mathsf{G}_3$
    $(\boldsymbol{s}_1, \boldsymbol{s}_2) \xleftarrow{\$} \mathcal{D}_{\boldsymbol{\Lambda}(\boldsymbol{B}),s,(\boldsymbol{c},\boldsymbol{0})}$                  // $\mathsf{G}_4$
until $\|(\boldsymbol{s}_1, \boldsymbol{s}_2)\|_2 \leq \beta$
$\sigma := (r, \boldsymbol{s}_2)$
$\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(m, \sigma)\}$
return $\sigma$

Oracle $\mathsf{H}(pk, r, m)$
if $\exists\, \boldsymbol{c} : (\boldsymbol{c}, pk, r, m) \in \mathcal{H}$
    return $\boldsymbol{c}$
$\boldsymbol{c} \xleftarrow{\$} \mathcal{R}_q$
$\mathcal{H} \leftarrow \mathcal{H} \cup \{(\boldsymbol{c}, pk, r, m)\}$
return $\boldsymbol{c}$

Oracle $\mathsf{H}'(\boldsymbol{h}, r, m)$
if $\exists\, \boldsymbol{c} : (\boldsymbol{c}, \boldsymbol{h}, r, m) \in \mathcal{H}$
    abort
$\boldsymbol{c} \xleftarrow{\$} \mathcal{R}_q$
$(\boldsymbol{s}_1, \boldsymbol{s}_2) := (\perp, \perp)$
$\boldsymbol{s}_1, \boldsymbol{s}_2 \leftarrow \mathcal{D}_{\mathcal{R},s}$                          // $\mathsf{G}_3 - \mathsf{G}_5$
$\boldsymbol{c} := \boldsymbol{s}_1 + \boldsymbol{s}_2 * \boldsymbol{h} \mod q$                // $\mathsf{G}_3 - \mathsf{G}_5$
$\mathcal{H} := \mathcal{H} \cup \{(\boldsymbol{c}, \boldsymbol{h}, r, m)\}$
return $(\boldsymbol{c}, \boldsymbol{s}_1, \boldsymbol{s}_2)$
```

**Figure 5.** Games for the proof of Theorem 2.

*Game* $\mathsf{G}_1$. 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 $p_{\mathsf{PreSmp},\beta} := \max_{\boldsymbol{c}\in\mathcal{R}_q} \Pr_{(\boldsymbol{s}_1,\boldsymbol{s}_2)\xleftarrow{\$}\mathsf{PreSmp}(\boldsymbol{B},s,(\boldsymbol{c},\boldsymbol{0}))}[\|(\boldsymbol{s}_1,\boldsymbol{s}_2)\|_2 \leq \beta]$ it holds that $\left|\Pr\left[\mathsf{G}_0^{\mathsf{A}} \Rightarrow 1\right] - \Pr\left[\mathsf{G}_1^{\mathsf{A}} \Rightarrow 1\right]\right| \leq \sum_{i=0}^{Q_s} \binom{C_s}{i} (1 - p_{\mathsf{PreSmp},\beta})^{C_s - i} (p_{\mathsf{PreSmp},\beta})^i$.

*Proof.* To proof the claim, we model the experiment using a binomial distributed random variable $X \sim B(C_s, p_{\mathsf{PreSmp},\beta})$, i.e. we have $C_s$ Bernoulli trials and success probability $p_{\mathsf{PreSmp},\beta}$. A trial corresponds to sampling a preimage using $\mathsf{PreSmp}$ in the signing oracle and the trial succeeds if the norm is sufficiently small, i.e. $\|(\boldsymbol{s}_1, \boldsymbol{s}_2)\|_2 \leq \beta$. 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* $\mathsf{G}_2$*.* This game is identical to the previous one except that it aborts in the signing oracle $\mathtt{Sgn}$ if there already exists a query to the random oracle on the same public key, salt $r$, and message $m$. To ease the depiction in further hybrids, we define a new RO $\mathsf{H}'$ maintaining the same set $\mathcal{H}$ as $\mathsf{H}$ but aborting in case of a query on the same input as a previous query. Oracle $\mathsf{H}'$ is then called within the signing oracle instead of $\mathsf{H}$.

Claim 2: $\left|\Pr\left[\mathsf{G}_1^{\mathsf{A}} \Rightarrow 1\right] - \Pr\left[\mathsf{G}_2^{\mathsf{A}} \Rightarrow 1\right]\right| \leq \frac{Q_s(C_s + Q_{\mathsf{H}})}{2^k}$.

*Proof.* The salt $r$ is chosen uniformly at random from $\{0,1\}^k$ for each signing query. The total number of elements in $\mathcal{H}$ is at most $C_s + Q_{\mathsf{H}}$, as at most one element is added per query to $\mathsf{H}/\mathsf{H}'$; and there are $C_s$ implicit queries via $\mathtt{Sgn}$ and $Q_{\mathsf{H}}$ direct ones. Thus, the probability that the freshly chosen salt was part of a previous query is at most $\frac{C_s + Q_{\mathsf{H}}}{2^k}$. For $Q_s$ queries to the signing oracle $\mathtt{Sgn}$, we obtain the claimed bound. ■

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

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

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

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

```
$\bar{\mathcal{P}}$
$(\boldsymbol{B}, \boldsymbol{h}) \xleftarrow{\$} \mathsf{TpdGen}$
for $i \in [Q]$
  $\boldsymbol{c}_i \xleftarrow{\$} \mathcal{R}_q$
return $((\boldsymbol{B}, \boldsymbol{h}), \boldsymbol{c}_1, \ldots, \boldsymbol{c}_Q)$
```

```
$\bar{\mathcal{Q}}$
$(\boldsymbol{B}, \boldsymbol{h}) \xleftarrow{\$} \mathsf{TpdGen}$
for $i \in [Q]$
  $(\boldsymbol{s}_1, \boldsymbol{s}_2) \leftarrow \mathcal{D}_{\mathcal{R},s}$
  $\boldsymbol{c}_i := \boldsymbol{s}_1 + \boldsymbol{s}_2 * \boldsymbol{h} \mod q$
return $((\boldsymbol{B}, \boldsymbol{h}), \boldsymbol{c}_1, \ldots, \boldsymbol{c}_Q)$
```

These distributions describe the underlying distributions of $\mathsf{G}_2$ and $\mathsf{G}_3$. By the data processing inequality (Lemma 9) it holds that, for any $a \in (1, \infty)$,

$$R_a(\mathsf{G}_2 \,||\, \mathsf{G}_3) \leq R_a(\bar{\mathcal{P}} \,||\, \bar{\mathcal{Q}}). \tag{1}$$

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

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

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

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

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

By Lemma 10 it follows

$$R_a(\bar{\mathcal{P}} \mid\mid \bar{\mathcal{Q}}) \leq \max_{(\cdot,\boldsymbol{h})\in\mathsf{TpdGen}} R_a(\mathcal{P} \mid\mid \mathcal{Q}_h)^Q. \tag{2}$$

Combining probability preservation (Lemma 8) with Equation (1) and Equation (2), we obtain

$$\Pr[\mathsf{G}_3^\mathsf{A} \Rightarrow 1] \geq \frac{\Pr[\mathsf{G}_2^\mathsf{A} \Rightarrow 1]^{\frac{a}{a-1}}}{R_a(\mathsf{G}_2 \mid\mid \mathsf{G}_3)} \geq \frac{\Pr[\mathsf{G}_2^\mathsf{A} \Rightarrow 1]^{\frac{a}{a-1}}}{\max_{(\cdot,\boldsymbol{h})\in\mathsf{TpdGen}} R_a(\mathcal{P} \mid\mid \mathcal{Q}_{\boldsymbol{h}})^Q}.$$

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

*Game* $\mathsf{G}_4$. This game is identical to the previous one except that the output of the preimage sampler $\mathsf{PreSmp}(\boldsymbol{B}, s, (\boldsymbol{c}, \boldsymbol{0}))$ is replaced by a Gaussian over the same lattice, standard deviation and center, namely $\mathcal{D}_{\boldsymbol{\Lambda}(\boldsymbol{B}),s,(\boldsymbol{c},\boldsymbol{0})}$.

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

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

*Proof.* The claim follows analogously to Game $\mathsf{G}_3$. ■

*Game* $\mathsf{G}_5$. This game is identical to the previous one except that preimages $\boldsymbol{s}_1, \boldsymbol{s}_2$ are not sampled from a Gaussian distribution centred at $(\boldsymbol{c}, \boldsymbol{0})$ as before. Instead, the preimages of $\boldsymbol{c}$ that were sampled in $\mathsf{H}'$ or $\mathsf{H}$ are used.

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

*Proof.* We need to show that the distributions of the games are the same. The RO output $\boldsymbol{c}$ is the same in both games. In $\mathsf{G}_4$, the signing oracle outputs $(\boldsymbol{s}_1, \boldsymbol{s}_2) \sim \mathcal{D}_{\boldsymbol{\Lambda}(\boldsymbol{B}),s,(\boldsymbol{c},\boldsymbol{0})}$. Since $\boldsymbol{\Lambda}(\boldsymbol{B})$ is the NTRU lattice for $\boldsymbol{h}$ and $q$ and the distribution is shifted by $(\boldsymbol{c}, \boldsymbol{0})$ the output is distributed according to a Gaussian $\mathcal{D}_{\mathcal{R},s}$ conditioned on $\boldsymbol{s}_1 + \boldsymbol{s}_2 * \boldsymbol{h} = \boldsymbol{c} \mod q$. The output distribution in Game $\mathsf{G}_5$ is a Gaussian $\mathcal{D}_{\mathcal{R},s}$ as well where the condition $\boldsymbol{s}_1 + \boldsymbol{s}_2 * \boldsymbol{h} = \boldsymbol{c} \mod q$ is fulfilled by construction (Line 29 and Line 46). ■

*Reduction from $\mathcal{R}$-**ISIS**.* Claim 6: There exists an adversary $\mathsf{B}$ against $\mathcal{R}$-**ISIS** such that

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

*Proof.* Adversary $\mathsf{B}$ is formally constructed in Figure 6. Due to the changes in the previous games, adversary $\mathsf{B}$ can perfectly simulate the game for adversary $\mathsf{A}$ against $\mathsf{G}_5$ without having the secret key for $\boldsymbol{h}$. Further they embed their own targets in the queries to $\mathsf{H}$. Let us assume, that $\mathsf{A}$ wins $\mathsf{G}_5$, i.e. the forgery verifies and $(m^\star, \cdot)$ was not queried before. This implies that there exists an $i^\star$ such that $\hat{\boldsymbol{c}}_{i^\star} = \boldsymbol{c}^\star$ because if $\mathsf{A}$ wins the game, the challenge RO output $\boldsymbol{c}^\star$ equals one of $\mathsf{B}$'s targets (that is exactly $\hat{\boldsymbol{c}}_{i^\star}$) or to a signing query. If it corresponds to a signing query, there is no way that adversary $\mathsf{A}$ can win the game due to the freshness condition $(m^\star, \cdot) \notin \mathcal{Q}$. Hence, Line 06 ensures the first winning condition of $\mathsf{B}$ which is $\boldsymbol{s}_1^\star + \boldsymbol{s}_2^\star * \boldsymbol{h} = \hat{\boldsymbol{c}} \mod q$. Further, the norm bound from $\mathsf{A}$ directly translates to the second winning condition, i.e. $\|(\boldsymbol{s}_1^\star, \boldsymbol{s}_2^\star)\|_2 \leq \beta$. ■

```
$\mathsf{B}(\boldsymbol{h}, \hat{\boldsymbol{c}}_1, \ldots, \hat{\boldsymbol{c}}_{Q_\mathsf{H}+1})$
$\mathcal{H}, \mathcal{Q} \leftarrow \emptyset$
$cnt, \ell \coloneqq 0$
$(m^\star, \sigma^\star) \xleftarrow{\$} \mathsf{A}^{\mathtt{Sgn}(\cdot), \mathsf{H}(\cdot,\cdot,\cdot)}(\boldsymbol{h})$
**parse** $\sigma^\star \to (r^\star, \boldsymbol{s}_2^\star)$
$\boldsymbol{c}^\star \coloneqq \mathsf{H}(\boldsymbol{h}, r^\star, m^\star)$
$\boldsymbol{s}_1^\star \coloneqq \boldsymbol{c}^\star - \boldsymbol{s}_2^\star * \boldsymbol{h} \mod q$
**find** $i^\star : \boldsymbol{c}^\star = \hat{\boldsymbol{c}}_{i^\star}$
**return** $(i^\star, \boldsymbol{s}_1^\star, \boldsymbol{s}_2^\star)$

**Oracle** $\mathtt{Sgn}(m)$
**return** $\mathsf{G}_5.\mathtt{Sgn}(m)$

**Oracle** $\mathsf{H}(pk, r, m)$
**if** $\exists\, \boldsymbol{c} : (\boldsymbol{c}, pk, r, m) \in \mathcal{H}$
    **return** $\boldsymbol{c}$
$\ell \coloneqq \ell + 1$
$\boldsymbol{c} \coloneqq \hat{\boldsymbol{c}}_\ell$        // embed challenge target
$\mathcal{H} \leftarrow \mathcal{H} \cup \{(\boldsymbol{c}, pk, r, m)\}$
**return** $\boldsymbol{c}$

**Oracle** $\mathsf{H}'(\boldsymbol{h}, r, m)$
**return** $\mathsf{G}_5.\mathsf{H}'(\boldsymbol{h}, r, m)$
```

**Figure 6.** Adversary B against $t$-$\mathcal{R}$-**ISIS** for the proof of Theorem 2.

## 6 Parameters and Analysing the Security Bound

In this section, we analyse the concrete security bounds for Falcon$^+$-512 and Falcon$^+$-1024 from Section 4.1. Recall that Falcon$^+$-512 and Falcon$^+$-1024 are slight modifications of Falcon-512 and Falcon-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 **SIS**/**ISIS** bit security, accounting for the specific norm bound from the theorem. Next, we analyse the bound in Theorem 1, beginning with proof constant $C_s$, denoting the maximal repetitions in the signing oracle. Then, we choose an $L$ which is close to $C_s$ such the additional loss in the following application of the Rényi argument is small. Next, based on the bit security of the **(I)SIS** 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 5 and Table 6, followed by a discussion of the findings.

### 6.1 Security of $\mathcal{R}$-SIS and $t$-$\mathcal{R}$-ISIS

We estimate the security of the $\mathcal{R}$-**SIS** and $t$-$\mathcal{R}$-**ISIS** terms in our bounds. We consider the $\mathcal{R}$-**SIS** and $t$-$\mathcal{R}$-**ISIS** problems (as defined in Definition 10), parametrised by a trapdoor generation algorithm TpdGen with trapdoor quality $\alpha$.

For strong unforgeability, Theorem 1 gives a reduction to $\mathcal{R}$-**SIS** with a norm bound of $2\beta$. For plain unforgeability Theorem 2 provides a reduction to $t$-$\mathcal{R}$-**ISIS** with a norm bound of $\beta$. For the hardness of $\mathcal{R}$-**SIS** / $t$-$\mathcal{R}$-**ISIS** 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 3).

We make the assumption that $\mathcal{R}$-**SIS** and $t$-$\mathcal{R}$-**ISIS** instances are as hard as random **SIS** and **ISIS** instances. Although it's possible that there are more efficient attacks against $t$-$\mathcal{R}$-**ISIS** [Ber22], we argue that a direct proof of $t$-$\mathcal{R}$-**ISIS** in Theorem 2 is meaningful, as it most accurately captures the security of the scheme. That is, the *plain unforgeability* not only requires the hardness of $t$-$\mathcal{R}$-**ISIS**, but is, in fact, also necessary. Specifically, an attack on $t$-$\mathcal{R}$-**ISIS** would directly imply an attack on 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*

| Assumption | Bit security |
|---|---|
| $t$-$\mathcal{R}$-**ISIS**$_{B=\beta_I}$ | 120 |
| $\mathcal{R}$-**SIS**$_{B=2\beta_I}$ | 95 |
| $t$-$\mathcal{R}$-**ISIS**$_{B=\beta_V}$ | 278 |
| $\mathcal{R}$-**SIS**$_{B=2\beta_V}$ | 0 |

**Table 4.** Bit security (Core-SVP) of the relevant $\mathcal{R}$-**ISIS**/$\mathcal{R}$-**SIS** instances with norm bound $B$.

*methodology"* from [ADPS16]. The resulting levels of bit security are summarised in Table 4. We refer to Figure 9 in Appendix C for the concrete prompts of the lattice estimator.[6]

### 6.2 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 $\lambda$, 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 17 (Optimal $C_s$).** For Falcon$^+$-512 with $\lambda = 128$ it holds that,

$$\underset{C_s}{\arg\min} \left\{ C_s \;\middle|\; \sum_{i=0}^{Q_s} \binom{C_s}{i} (1 - p_{\mathsf{PreSmp},\beta})^{C_s - i} (p_{\mathsf{PreSmp},\beta})^i \leq 2^\lambda \right\} \leq 2^{64} + 2^{50},$$

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

$$\underset{C_s}{\arg\min} \left\{ C_s \;\middle|\; \sum_{i=0}^{Q_s} \binom{C_s}{i} (1 - p_{\mathsf{PreSmp},\beta})^{C_s - i} (p_{\mathsf{PreSmp},\beta})^i \leq 2^\lambda \right\} \leq 2^{64} + 2^{36}.$$

The proof can be found in Appendix C.1. For different values $Q_s$, $C_s$ can be computed in the same way, as shown in Tables 5 and 6.

### 6.3 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 16 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 2), 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.

**Corollary 3 (Rényi Loss for Falcon$^+$-512 (Preimage Sampler) in Thm. 2).** For $\varepsilon \geq 2^{-\lambda} = 2^{-120}$, $r_p = R_{a_p}(\mathsf{PreSmp} \;||\; \mathcal{D})$, $C_s = 2^{64} + 2^{50}$, 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 $a_p \approx 72.96$.

[6] The lattice estimator results in 228 bits of security for Falcon$^+$-1024 when $n = 1024$ and $B = 2\beta$. This is clearly incorrect, as $2\beta = 16764.88 > q = 12289$.

The proof can be found in Appendix C.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. 2).** For $\varepsilon \geq 2^{-\lambda} = 2^{-116.5}$, $r_u = R_{a_u}(\mathcal{U}(\mathcal{R}_q) \,||\, \mathcal{U}_{\boldsymbol{h}})$, $C_s = 2^{64} + 2^{50}$, 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 Appendix C.3.

FALCON$^+$-1024. We apply the same arguments as for FALCON$^+$-512. 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. 2).** For $\varepsilon \geq 2^{-\lambda} = 2^{-278}$, $r_p = R_{a_p}(\mathsf{PreSmp} \,||\, \mathcal{D})$, $C_s = 2^{64} + 2^{36}$, 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. 2).** For $\varepsilon \geq 2^{-\lambda} = 2^{-274}$, $r_u = R_{a_u}(\mathcal{U}(\mathcal{R}_q) \,||\, \mathcal{U}_{\boldsymbol{h}})$, $C_s = 2^{64} + 2^{36}$, 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. ■

OTHER BOUNDS AND NUMBER OF SIGNING QUERIES. The optimal Rényi orders for the strong unforgeability bound (Theorem 1) 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.

### 6.4 Final Security and Discussion

To conclude the analysis of the bounds, we note that the term $Q_s(C_s{+}Q_{\mathsf{H}})/2^k$ provides $\lambda$ bits of security when $k \geq \log(Q_s) + \lambda$. For both parameter sets, FALCON$^+$ achieves this by selecting $k = 320$. The binomial term fulfils $\lambda$ bits of security by choosing an appropriate $C_s$, as detailed in the proof of Lemma 17. An overview of the results from the previous subsections is presented in Table 5 for FALCON$^+$-512 and in Table 6 for 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. Although the bit security for plain unforgeability is close to the target, the bit security for strong unforgeability in FALCON$^+$-512 is insufficient. Specifically, for $Q_s = 2^{64}$, the security level is only 89 bits, offering no meaningful security guarantee. For FALCON$^+$-1024, the situation is even worse, as

[7] 270 bits refers to the bit security of the computational term. See Section 6.4 for more information.

**NIST Level I:** $\mathcal{R} = \mathbb{Z}_{12289}[X]/(X^{512}+1)$

| Notion<br>Parameter | **SUF-CMA**<br>Thm. 1 | | **UF-CMA**<br>Thm. 2 | |
|---|---|---|---|---|
| $t$-$\mathcal{R}$-**ISIS**/$\mathcal{R}$-**SIS** length bound $B$ | $2\beta = 11667.86$ | | $\beta = 5833.93$ | |
| Bit security (core-SVP), $t$-$\mathcal{R}$-$\mathbf{ISIS}_{m=1,q=q,\alpha=1.17,B=\beta}$ | — | | 120 | |
| Bit security (core-SVP), $\mathcal{R}$-$\mathbf{SIS}_{m=1,q=q,\alpha=1.17,B=2\beta}$ | 95 | | — | |
| Max Signing queries, $Q_s$ | $2^{58}$ | $2^{64}$ | $2^{58}$ | $2^{64}$ |
| Max repetitions, $C_s(\lambda, Q_s)$ | $2^{58}+2^{44}$ | $2^{64}+2^{50}$ | $2^{58}+2^{44}$ | $2^{64}+2^{50}$ |
| Number of programmed RO queries, $L$ | $2^{58}$ | $2^{60}$ | — | |
| Rényi Order, $a_p$ | 519.33 | 64.92 | 583.67 | 72.96 |
| Rényi Order, $a_u$ | 366.26 | 61.98 | 582.46 | 71.73 |
| Bits lost from Rényi $a_p$ | 0.5 | 3 | 0.5 | 3.5 |
| Bits lost from Rényi $a_u$ | 0.5 | 3 | 0.5 | 3.5 |
| **Final bit security** | **94** | **89** | **119** | **113** |

**Table 5.** Provable security level of Falcon$^+$-512.

**NIST Level V:** $\mathcal{R} = \mathbb{Z}_{12289}[X]/(X^{1024}+1)$

| Notion<br>Parameter | **SUF-CMA**<br>Thm. 1 | **UF-CMA**<br>Thm. 2 |
|---|---|---|
| $t$-$\mathcal{R}$-**ISIS**/$\mathcal{R}$-**SIS** length bound $B$ | $2\beta = 16764.87$ | $\beta = 8382.44$ |
| Bit security (core-SVP), $t$-$\mathcal{R}$-$\mathbf{ISIS}_{m=1,q=q,\alpha=1.17,B=\beta}$ | — | 278 |
| Bit security (core-SVP), $\mathcal{R}$-$\mathbf{SIS}_{m=1,q=q,\alpha=1.17,B=2\beta}$ | 0 | — |
| Max Signing queries, $Q_s$ | n/a | $2^{64}$ |
| Max repetitions, $C_s(\lambda, Q_s)$ | | $2^{64}+2^{36}$ |
| Rényi Order, $a_p$ | | 157.05 |
| Rényi Order, $a_u$ | | 155.92 |
| Bits lost from Rényi $a_p$ | | 4 |
| Bits lost from Rényi $a_u$ | | 4 |
| **Final bit security** | **0** | **256 (270)**[7] |

**Table 6.** Provable security level of Falcon$^+$-1024.

no security can be proven due to the norm bound $2\beta$ exceeding the modulus $q$. If strong unforgeability could be reduced to one-wayness (or **ISIS**), as in the case of plain unforgeability, rather than collision resistance (or **SIS**), a smaller norm bound and better security could be achieved. However, current proof techniques cannot address this problem, and we believe it is unlikely to be feasible in general. Achieving strong unforgeability for Falcon$^+$ would thus require increasing both the ring dimension and modulus, leading to larger public keys and signature sizes.

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. This is necessary because, the Falcon$^+$ parameters do not account for additional queries caused by the signing procedure's repetition. Allowing $2^{64}$ queries increases the Rényi loss, which is problematic, given that the security of **ISIS** is already tightly set to meet the target level. For Falcon$^+$-1024, the larger security margin between **ISIS** and the target security compensates for higher Rényi losses, making this issue less critical. Therefore, we also present the maximum number of signing queries that can be supported while maintaining a Rényi loss of at most 1 bit. This issues is not an artifact of our proof strategy but arises from the inherent repetition in the signing procedure and the sensitivity of the Rényi arguments. To support $2^{64}$ signing queries with tight Rényi bounds, the smoothing parameter error $\epsilon$ would need to account for the maximum repetitions, leading to $\epsilon = (C_s \cdot \lambda)^{-1/2}$. This would, however, increase parameters such as the signature size.

**Acknowledgements.** The authors thank Doreen Riepel for suggesting the application of [BRTZ24] to improve the security bounds, and Melissa Chase for pointing out how to support more signing queries. The authors also thank Mihir Bellare, Joel Gärtner, Thomas Prest, and Damien Stehlé for helpful feedback and comments. Part of the work was done while Phillip Gajland was at IBM Research – Zurich. Phillip Gajland was supported by the 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. (Cited on page 21.)

[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 10.)

[Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *28th Annual ACM Symposium on Theory of Computing*, pages 99–108, Philadephia, 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 7 and 20.)

[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 7 and 20.)

[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 10.)

[BBD+23] Joppe W. Bos, Olivier Bronchain, Léo Ducas, Serge Fehr, Yu-Hsuan Huang, Thomas Pornin, Eamonn W. Postlethwaite, Thomas Prest, Ludo N. Pulles, and Wessel van Woerden. HAWK. 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 7.)

[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 12 and 20.)

[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, Tancrède Lepoint, Damien Stehlé, and Ron Steinfeld. Improved security proofs in lattice-based cryptography: Using the Rényi divergence rather than the statistical distance. In Tetsu Iwata and Jung Hee Cheon, editors, *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, 11, and 13.)

[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 page 5.)

[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.)

[BR04] Mihir Bellare and Phillip Rogaway. Code-based game-playing proofs and the security of triple encryption. Cryptology ePrint Archive, Report 2004/331, 2004. URL: https://eprint.iacr.org/2004/331. (Cited on page 8.)

[BRTZ24] Mihir Bellare, Doreen Riepel, Stefano Tessaro, and Yizhao Zhang. Count corruptions, not users: Improved tightness for signatures, encryption and authenticated key exchange. In Kai-Min Chung and Yu Sasaki, editors, *Advances in Cryptology – ASIACRYPT 2024, Part II*, volume 15485 of *Lecture Notes in Computer Science*, pages 326–360, Kolkata, India, December 9–13, 2024. Springer, Singapore, Singapore. doi:10.1007/978-981-96-0888-1_11. (Cited on pages 6, 15, and 24.)

[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 14.)

[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 29.)

[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 14.)

[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, 5, and 10.)

[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 10.)

[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, 13, and 14.)

[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, 13, and 14.)

[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.)

[FH23] Serge Fehr and Yu-Hsuan Huang. On the quantum security of HAWK. In Thomas Johansson and Daniel Smith-Tone, editors, *Post-Quantum Cryptography - 14th International Workshop, PQCrypto 2023*, pages 405–416, College Park, USA, August 16–18, 2023. Springer, Cham, Switzerland. doi:10.1007/978-3-031-40003-2_15. (Cited on page 7.)

[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 13.)

[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, 10, 12, 13, 29, and 30.)

[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 5.)

[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 11.)

[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, 5, and 10.)

[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, 7, 12, and 15.)

[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, Tancrède Lepoint, Peter Schwabe, Gregor Seiler, Damien Stehlé, and Shi Bai. CRYSTALS-DILITHIUM. Technical report, National Institute of Standards and Technology, 2022. available at https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022. (Cited on pages 3 and 4.)

[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 12.)

[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 11.)

[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 10.)

[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 page 10.)

[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 page 10.)

[MW17] Daniele Micciancio and Michael Walter. Gaussian sampling over the integers: Efficient, generic, constant-time. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology – CRYPTO 2017, Part II*, volume 10402 of *Lecture Notes in Computer Science*, pages 455–485, Santa Barbara, CA, USA, August 20–24, 2017. Springer, Cham, Switzerland. doi:10.1007/978-3-319-63715-0_16. (Cited on page 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 5, 15, and 16.)

[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: 1cdc9c5bdd5b5894475febd7e23abbcb5056197b. (Cited on pages 4 and 6.)

[PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In Shai Halevi and Tal Rabin, editors, *TCC 2006: 3rd Theory of Cryptography Conference*, volume 3876 of *Lecture Notes in Computer Science*, pages 145–166, New York, NY, USA, March 4–7, 2006. Springer Berlin Heidelberg, Germany. doi:10.1007/11681878_8. (Cited on pages 11 and 29.)

[Pre15] Thomas Prest. *Gaussian sampling in lattice-based cryptography.* PhD thesis, Ecole normale supérieure-ENS PARIS, 2015. (Cited on pages 10 and 12.)

[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 6, 7, 11, 12, 15, and 21.)

[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 14.)

[Rén61] Alfréd Rényi. On measures of entropy and information. Proc. 4th Berkeley Symp. Math. Stat. Probab. 1, 547-561 (1961)., 1961. (Cited on page 11.)

[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, Tancrède Lepoint, Vadim Lyubashevsky, John M. Schanck, Gregor Seiler, Damien Stehlé, and Jintai Ding. CRYSTALS-KYBER. Technical report, National Institute of Standards and Technology, 2022. available at https://csrc.nist.gov/Projects/post-quantum-cryptography/selected-algorithms-2022. (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, 14, and 21.)

[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 11.)

[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 13.)

# A Proofs for Section 2 and Section 3

## A.1 Proof of Lemma 6

**Lemma 6 (Min-Entropy of Gaussian (implicit in [PR06, Lem. 2.10])).** Let $n \in \mathbb{N}$, $\mathbf{\Lambda} = \mathbb{Z}^n$, $\epsilon \in (0, \frac{1}{2})$ and $s \geq \eta_\epsilon(\mathbf{\Lambda}) > 2$. Then for any $c \in \mathbb{R}^n$ and $x \in \mathbf{\Lambda}$,

$$\mathcal{D}_{\mathbf{\Lambda},s,c}(x) \leq \frac{1}{2^{n-1}}.$$

*Proof.* Note that $\det(\mathbf{\Lambda}) = \det(\mathbb{Z}^n) = 1$. Since $\rho_{s,c}(x) \leq 1$ and by applying Lemma 5,

$$\mathcal{D}_{\mathbf{\Lambda},s,c}(x) = \frac{\rho_{s,c}(x)}{\rho_{s,c}(\mathbf{\Lambda})} \leq \frac{1}{(1-\epsilon) \cdot \frac{s^n}{\det(\mathbf{\Lambda})}} = \frac{1}{s^n(1-\epsilon)}.$$

Finally, since $\epsilon < \frac{1}{2}$ and $s > 2$, we get $\frac{1}{s^n(1-\epsilon)} < \frac{1}{2^n(1/2)} = \frac{1}{2^{n-1}}$. ■

## A.2 Proof of Lemma 13

**Lemma 13 (Rényi Divergence of Gaussian Sample over $\mathbf{\Lambda}/\mathbf{\Lambda}'$ (adapted from [GPV08, Cor. 2.8])).** Let $\mathbf{\Lambda}, \mathbf{\Lambda}'$ be $n$-dimensional full-rank lattices with $\mathbf{\Lambda}' \subseteq \mathbf{\Lambda}$. Then for any $a \in (1, \infty)$, $\epsilon \in (0, \frac{1}{2})$, any $s \geq \eta_\epsilon(\mathbf{\Lambda}')$, and any $c \in \mathbb{R}^n$,

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

*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 $\mathbf{\Lambda}/\mathbf{\Lambda}'$ is defined as the additive group of cosets $x + \mathbf{\Lambda}', x \in \mathbf{\Lambda}$. Sampling from a discrete Gaussian over this quotient group we obtain that for any $x \in \mathbf{\Lambda}$

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

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

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

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

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

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 11 with $\delta = \frac{2\epsilon}{1-\epsilon}$, we obtain

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

This completes the proof. ■

### A.3 Proof of Lemma 14

**Lemma 14 (Rényi divergence (adapted from [GPV08, Lem 5.2])).** If the columns of $\boldsymbol{A} \in \mathbb{Z}_q^{n \times m}$ generate $\mathbb{Z}_q^n$, $a \in (1, \infty)$, $\epsilon \in (0, \frac{1}{2})$, and $s \geq \eta_\epsilon(\boldsymbol{\Lambda}^\perp(\boldsymbol{A}))$; 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 $u = \boldsymbol{A}e \mod q$, it holds that

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

*Proof.* For simplicity we denote $\boldsymbol{\Lambda}^\perp = \boldsymbol{\Lambda}^\perp(\boldsymbol{A})$. By assumption the set of all syndromes of $\boldsymbol{A}$ equals $\mathbb{Z}_q^n$, i.e. $\{\boldsymbol{A}e \mod q \mid e \in \mathbb{Z}^m\} = \mathbb{Z}_q^n$. Consider the quotient group $(\mathbb{Z}^m/\boldsymbol{\Lambda}^\perp)$ which is defined as the group of all cosets, i.e. $\{e + \boldsymbol{\Lambda}^\perp \mid e \in \mathbb{Z}^m\}$. This quotient group is isomorphic to the set of syndromes of $\boldsymbol{A}$ via the mapping $e + \boldsymbol{\Lambda}^\perp \mapsto \boldsymbol{A}e \mod q$, where $e \in \mathbb{Z}^m$. Hence, we have $\mathcal{P} \simeq \mathcal{U}(\mathbb{Z}^m/\boldsymbol{\Lambda}^\perp)$. Further, the distribution $\mathcal{D}_{\mathbb{Z}^m/\boldsymbol{\Lambda}^\perp,s} = \mathcal{D}_{\mathbb{Z}^m,s} \mod \boldsymbol{\Lambda}^\perp$ is the distribution of $e \sim \mathcal{D}_{\mathbb{Z}^m,s}$ reduced modulo $\boldsymbol{\Lambda}^\perp$. That is, the coset $e + \boldsymbol{\Lambda}^\perp$ 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 13 with $\boldsymbol{\Lambda} = \mathbb{Z}^m$, $\boldsymbol{\Lambda}' = \boldsymbol{\Lambda}^\perp$ and $c = 0$ to obtain the claim. ■

### A.4 Proof of Corollary 2

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

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

*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 $\boldsymbol{h} \in \mathcal{R}_q \setminus \{\boldsymbol{0}\}$, we consider matrix $\boldsymbol{A} = \left[I_N \middle| \mathcal{A}(\boldsymbol{h})\right] \in \mathbb{Z}^{n \times 2n}$ that defines the NTRU lattice $\boldsymbol{\Lambda}_{\boldsymbol{h},q} = \boldsymbol{\Lambda}^\perp(\boldsymbol{A})$. By Lemma 1 the anticircuclant 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 $(\boldsymbol{u}, \boldsymbol{v})$ with $\boldsymbol{u}, \boldsymbol{v} \sim \mathcal{D}_{\mathcal{R},s}$ as well as for the resulting distributions $\boldsymbol{A} \cdot \mathcal{A}(e) \mod q$ and the distribution of $\boldsymbol{z}$ such that $\mathcal{A}(\boldsymbol{z}) = \boldsymbol{A} \begin{bmatrix} \mathcal{A}(\boldsymbol{u}) \\ \mathcal{A}(\boldsymbol{v}) \end{bmatrix} = \mathcal{A}(\boldsymbol{u}) + \mathcal{A}(\boldsymbol{h}) \cdot \mathcal{A}(\boldsymbol{v}) \mod q$. The latter distribution is equivalent to $\mathcal{Q}$. Finally, due to its special structure with identity $I_N$ on the left, $\boldsymbol{A}$ generates $\mathbb{Z}_q^n$ such that we can apply Lemma 14 to conclude the proof. ■

### A.5 Proof of Lemma 15

**Lemma 15 (Relative Error for Tailbounds).** Let $\mathcal{P}$ and $\mathcal{Q}$ be two distributions with $\sup(\mathcal{P}) = \sup(\mathcal{Q}) = \mathbb{Z}^n$ such that their relative error is bounded by $\frac{\mathcal{P}}{\mathcal{Q}} \leq 1 + \delta$. 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 $\infty$:

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

Further, let $E$ be the event that the drawn value $x$ fulfils $\|x\|_2 > \beta$. Applying the probability preservation for $R_\infty$ (Lemma 8) we obtain

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

■

### A.6 Proof of Lemma 16

**Lemma 16 (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 $Q \in \mathbb{N}$ and $a \in (1, \infty)$ the Rényi divergence between two arbitrary distributions is at most $R_a \in [1, \infty)$, and

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

Then

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

*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 $\mathcal{E}_2$. ■

# B Proof of Theorem 1

**Theorem 1 (Strong Unforgeability).** For any adversary $\mathsf{A}$ against the **SUF-CMA** security of $\textsc{CoreFalcon}^{+}$ (Figure 4) running in time $t_{\mathsf{A}}$, making at most $Q_s$ signing queries and $Q_{\mathsf{H}}$ random oracle queries, there exist an adversary $\mathsf{B}$ against $\mathcal{R}$-**SIS** running in time $t_{\mathsf{B}} \approx t_{\mathsf{A}}$ such that for all $C_s \in \mathbb{N}^{\geq 1}, L \in [Q_{\mathsf{H}}+1]$ and $a_u, a_p \in \mathbb{R}^{>1}$ it holds

$$\mathrm{Adv}^{Q_s\text{-}\mathbf{SUF\text{-}CMA}}_{\textsc{CoreFalcon}^{+},\mathsf{A}} \leq$$

$$\frac{(Q_{\mathsf{H}}+1)}{L \cdot p_{\mathsf{PreSmp},\beta}} \cdot \left( r_u^{C_s+L} \cdot \left( r_p^{C_s+1} \cdot \left( \mathrm{Adv}^{\mathcal{R}\text{-}\mathbf{SIS}}_{1,q,\alpha,2\beta,\mathsf{B}} + \frac{1}{2^{n-1}} \right) \right)^{\frac{a_p-1}{a_p}} \right)^{\frac{a_u-1}{a_u}}$$

$$+ \sum_{i=0}^{Q_s} \binom{C_s}{i} (1-p_{\mathsf{PreSmp},\beta})^{C_s-i} (p_{\mathsf{PreSmp},\beta})^{i} + \frac{Q_s(C_s+Q_{\mathsf{H}})}{2^k} \ ,$$

where

$p_{\mathsf{PreSmp},\beta} := \max_{(\boldsymbol{B},\cdot)\in\mathsf{TpdGen},\boldsymbol{c}\in\mathcal{R}_q} \Pr_{(\boldsymbol{s}_1,\boldsymbol{s}_2)\xleftarrow{\$}\mathsf{PreSmp}(\boldsymbol{B},s,(\boldsymbol{c},\boldsymbol{0}))}[\|(\boldsymbol{s}_1,\boldsymbol{s}_2)\|_2 \leq \beta]$,
$r_u = \max_{(\cdot,\boldsymbol{h})\in\mathsf{TpdGen}} R_{a_u}(\mathcal{P} \mid\mid \mathcal{Q}_{\boldsymbol{h}})$ with $\mathcal{P} = \mathcal{U}(\mathcal{R}_q)$ and $\mathcal{Q}_{\boldsymbol{h}}$ the distribution of $\boldsymbol{s}_1 + \boldsymbol{s}_2 * \boldsymbol{h} \mod q$, where $\boldsymbol{s}_1, \boldsymbol{s}_2 \sim \mathcal{D}_{\mathcal{R},s}$,
$r_p = \max_{(\boldsymbol{B},\cdot)\in\mathsf{TpdGen},\boldsymbol{c}\in\mathcal{R}_q} R_{a_p}(\mathsf{PreSmp}(\boldsymbol{B},s,(\boldsymbol{c},\boldsymbol{0})) \mid\mid \mathcal{D}_{\Lambda(\boldsymbol{B},s,(\boldsymbol{c},\boldsymbol{0}))})$.

*Proof.* Consider the sequence of games depicted in Figure 7.

*Game* $\mathsf{G}_0$. This is the strong unforgeability game for $\textsc{CoreFalcon}^{+}$ so by definition we have

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

*Game* $\mathsf{G}_1$. 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 7: For $p_{\mathsf{PreSmp},\beta} := \max_{\boldsymbol{c}\in\mathcal{R}_q} \Pr_{(\boldsymbol{s}_1,\boldsymbol{s}_2)\xleftarrow{\$}\mathsf{PreSmp}(\boldsymbol{B},s,(\boldsymbol{c},\boldsymbol{0}))}[\|(\boldsymbol{s}_1,\boldsymbol{s}_2)\|_2 \leq \beta]$ it holds that $\left|\Pr\left[\mathsf{G}_0^{\mathsf{A}} \Rightarrow 1\right] - \Pr\left[\mathsf{G}_1^{\mathsf{A}} \Rightarrow 1\right]\right| \leq \sum_{i=0}^{Q_s} \binom{C_s}{i} (1-p_{\mathsf{PreSmp},\beta})^{C_s-i}(p_{\mathsf{PreSmp},\beta})^{i}$.

*Proof.* To proof the claim, we model the experiment using a binomial distributed random variable $X \sim B(C_s, p_{\mathsf{PreSmp},\beta})$, i.e. we have $C_s$ Bernoulli trials and success probability $p_{\mathsf{PreSmp},\beta}$. A trial corresponds to sampling a preimage using $\mathsf{PreSmp}$ in the signing oracle and the trial succeeds if the norm is sufficiently small, i.e. $\|(\boldsymbol{s}_1,\boldsymbol{s}_2)\|_2 \leq \beta$. 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* $\mathsf{G}_2$. This game is identical to the previous one except that it aborts in the signing oracle `Sgn` if there already exists a query to the random oracle on the same public key, salt $r$, and message $m$. To ease the depiction in further hybrids, we define a new RO $\mathsf{H}'$ maintaining the same set $\mathcal{H}$ as $\mathsf{H}$ but aborting in case of a query on the same input as a previous query. Oracle $\mathsf{H}'$ is then called within the signing oracle instead of $\mathsf{H}$.

Claim 8: $\left|\Pr\left[\mathsf{G}_1^{\mathsf{A}} \Rightarrow 1\right] - \Pr\left[\mathsf{G}_2^{\mathsf{A}} \Rightarrow 1\right]\right| \leq \frac{Q_s(C_s+Q_{\mathsf{H}})}{2^k}$.

*Proof.* The salt $r$ is chosen uniformly at random from $\{0,1\}^k$ for each signing query. The total number of elements in $\mathcal{H}$ is at most $C_s + Q_{\mathsf{H}}$, as at most one element is added per query to $\mathsf{H}/\mathsf{H}'$; and there are $C_s$ implicit queries via `Sgn` and $Q_{\mathsf{H}}$ direct ones. Thus, the probability that the freshly chosen salt was part of a previous query is at most $\frac{C_s+Q_{\mathsf{H}}}{2^k}$. For $Q_s$ queries to the signing oracle `Sgn`, we obtain the claimed bound. ■

```
Games G_0 − G_7
01 $\mathcal{H}, \mathcal{Q} \leftarrow \emptyset$
02 $cnt, \ell := 0$
03 choose random $\mathcal{S} \subseteq [Q_{\mathsf{H}}+1]$ such that $|\mathcal{S}| = L$    // $G_4 - G_7$
04 $(\boldsymbol{B}, \boldsymbol{h}) \xleftarrow{\$} \mathsf{Gen}$
05 $(m^\star, \sigma^\star) \xleftarrow{\$} \mathsf{A}^{\mathsf{Sgn}(\cdot), \mathsf{H}(\cdot,\cdot,\cdot)}(\boldsymbol{h})$
06 parse $\sigma^\star \rightarrow (r^\star, \boldsymbol{s}_2^\star)$
07 $\boldsymbol{c}^\star := \mathsf{H}(\boldsymbol{h}, r^\star, m^\star)$    // $G_3 - G_7$
08 $(\boldsymbol{s}_1', \boldsymbol{s}_2') \xleftarrow{\$} \mathsf{PreSmp}(\boldsymbol{B}, s, (\boldsymbol{c}^\star, \boldsymbol{0}))$    // $G_3 - G_7$
09 $(\boldsymbol{s}_1', \boldsymbol{s}_2') \xleftarrow{\$} \mathcal{D}_{\boldsymbol{\Lambda}(\boldsymbol{B}), s, (\boldsymbol{c}^\star, \boldsymbol{0})}$    // $G_6 - G_7$
10 find $(\boldsymbol{s}_1', \boldsymbol{s}_2') : (\boldsymbol{c}^\star, \boldsymbol{h}, r^\star, m^\star, \cdot, \boldsymbol{s}_1', \boldsymbol{s}_2') \in \mathcal{H}$    // $G_7 - G_7$
11 if $\|(\boldsymbol{s}_1', \boldsymbol{s}_2')\|_2 > \beta$    // $G_3 - G_7$
12   abort    // $G_3 - G_7$
13 find $\ell^\star : (\boldsymbol{c}^\star, \boldsymbol{h}, r^\star, m^\star, \ell^\star, \cdot, \cdot) \in \mathcal{H}$    // $G_4 - G_7$
14 if $\ell^\star \notin \mathcal{S} \wedge \ell^\star \neq \perp$    // $G_4 - G_7$
15   abort    // $G_4 - G_7$
16 $\boldsymbol{s}_1^\star := \boldsymbol{c}^\star - \boldsymbol{s}_2^\star * \boldsymbol{h} \mod q$
17 return $[\![\mathsf{Ver}(\boldsymbol{h}, m^\star, \sigma^\star) = 1 \wedge (m^\star, \sigma^\star) \notin \mathcal{Q}]\!]$
```

```
Oracle H'(h, r, m)
18 if $\exists\, \boldsymbol{c} : (\boldsymbol{c}, \boldsymbol{h}, r, m, \cdot, \cdot, \cdot) \in \mathcal{H}$
19   abort
20 $\boldsymbol{c} \xleftarrow{\$} \mathcal{R}_q$
21 $(\boldsymbol{s}_1, \boldsymbol{s}_2) := (\perp, \perp)$
22 $\boldsymbol{s}_1, \boldsymbol{s}_2 \leftarrow \mathcal{D}_{\mathcal{R},s}$    // $G_5 - G_7$
23 $\boldsymbol{c} := \boldsymbol{s}_1 + \boldsymbol{s}_2 * \boldsymbol{h} \mod q$    // $G_5 - G_7$
24 $\mathcal{H} := \mathcal{H} \cup \{(\boldsymbol{c}, \boldsymbol{h}, r, m, \perp, \boldsymbol{s}_1, \boldsymbol{s}_2)\}$
25 return $(\boldsymbol{c}, \boldsymbol{s}_1, \boldsymbol{s}_2)$
```

```
Oracle Sgn(m)
26 repeat
27   $cnt \leftarrow cnt + 1$    // $G_1 - G_7$
28   if $cnt > C_s$    // $G_1 - G_7$
29     abort    // $G_1 - G_7$
30   $r \xleftarrow{\$} \{0,1\}^k$
31   $\boldsymbol{c} := \mathsf{H}(\boldsymbol{h}, r, m)$    // $G_0 - G_1$
32   $(\boldsymbol{c}, \boldsymbol{s}_1, \boldsymbol{s}_2) := \mathsf{H}'(\boldsymbol{h}, r, m)$    // $G_2 - G_7$
33   $(\boldsymbol{s}_1, \boldsymbol{s}_2) \xleftarrow{\$} \mathsf{PreSmp}(\boldsymbol{B}, s, (\boldsymbol{c}, \boldsymbol{0}))$    // $G_0 - G_5$
34   $(\boldsymbol{s}_1, \boldsymbol{s}_2) \xleftarrow{\$} \mathcal{D}_{\boldsymbol{\Lambda}(\boldsymbol{B}), s, (\boldsymbol{c}, \boldsymbol{0})}$    // $G_6$
35 until $\|(\boldsymbol{s}_1, \boldsymbol{s}_2)\|_2 \leq \beta$
36 $\sigma := (r, \boldsymbol{s}_2)$
37 $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{(m, \sigma)\}$
38 return $\sigma$
```

```
Oracle H(pk, r, m)
39 if $\exists\, \boldsymbol{c} : (\boldsymbol{c}, pk, r, m, \cdot, \cdot, \cdot) \in \mathcal{H}$
40   return $\boldsymbol{c}$
41 $\boldsymbol{c} \xleftarrow{\$} \mathcal{R}_q$
42 $(\boldsymbol{s}_1, \boldsymbol{s}_2) := (\perp, \perp)$
43 $\ell := \ell + 1$
44 if $\ell \in \mathcal{S}$    // $G_5 - G_7$
45   $\boldsymbol{s}_1, \boldsymbol{s}_2 \leftarrow \mathcal{D}_{\mathcal{R},s}$    // $G_5 - G_7$
46   $\boldsymbol{c} := \boldsymbol{s}_1 + \boldsymbol{s}_2 * \boldsymbol{h} \mod q$    // $G_5 - G_7$
47 $\mathcal{H} \leftarrow \mathcal{H} \cup \{(\boldsymbol{c}, pk, r, m, \ell, \boldsymbol{s}_1, \boldsymbol{s}_2)\}$
48 return $\boldsymbol{c}$
```

**Figure 7.** Games for the proof of Theorem 1.

*Game* $\mathsf{G}_3$. This game is identical to the previous one except that the game computes a preimage of $\boldsymbol{c}^\star$ with respect to $\boldsymbol{h}$ using the preimage sampler PreSmp with trapdoor $\boldsymbol{B}$. If the norm of the resulting preimage $(\boldsymbol{s}_1', \boldsymbol{s}_2')$ is larger than $\beta$, the game aborts.

Claim 9: $\Pr[\mathsf{G}_2^{\mathsf{A}} \Rightarrow 1] = (p_{\mathsf{PreSmp},\beta})^{-1} \cdot \Pr[\mathsf{G}_3^{\mathsf{A}} \Rightarrow 1]$.

*Proof.* The probability that the abort event does not occur is $p_{\mathsf{PreSmp},\beta}$. Since the preimage is computed after adversary A outputs their forgery and the winning condition is not affected by $(\boldsymbol{s}_1', \boldsymbol{s}_2')$, the abort event is independent of A's winning probability. Hence we can apply the multiplicative difference lemma to obtain the statement. ■

*Game* $\mathsf{G}_4$. This game is identical to the previous one except for the following changes. For a fixed integer $L \in [Q_{\mathsf{H}}+1]$, we choose a random subset of $[Q_{\mathsf{H}}+1]$ of size $L$ in the beginning of the game which we call $\mathcal{S}$. For every query to H, we record the query number, $\ell$, and the game aborts if the the random oracle query of the forgery, Line 07, corresponds to a query to H and the query number, $\ell^\star$ is not in $\mathcal{S}$.

Claim 10: $\Pr[\mathsf{G}_3^{\mathsf{A}} \Rightarrow 1] \leq \frac{Q_{\mathsf{H}}+1}{L} \cdot \Pr[\mathsf{G}_4^{\mathsf{A}} \Rightarrow 1]$.

*Proof.* First note that the probability of an abort is independent of the winning probability in $\mathsf{G}_3$ since set $\mathcal{S}$ is not used anywhere else and hence does not influence the winning probability. The game does not abort if the challenge RO query corresponds to a signing query (case $\ell^\star = \bot$) but only to a direct RO query. Since $\mathcal{S}$ is of fixed size $L$, we can apply the multiplicative difference lemma and obtain the claim. ■

*Game* $\mathsf{G}_5$. This game is the same as the previous one, except that random oracle $\mathsf{H}'$ no longer returns a uniformly random element $\boldsymbol{c} \xleftarrow{\$} \mathcal{R}_q$. Instead, it computes $\boldsymbol{c}$ as follows: It samples two elements $\boldsymbol{s}_1, \boldsymbol{s}_2$ from a Gaussian distribution $\mathcal{D}_{\mathcal{R},s}$ with standard deviation $s$ over ring $\mathcal{R}$. Then, $\boldsymbol{c}$ is computed as $\boldsymbol{c} = \boldsymbol{s}_1 + \boldsymbol{s}_2 * \boldsymbol{h} \mod q$, where $\boldsymbol{h}$ is the public key. For future use, $\boldsymbol{s}_1, \boldsymbol{s}_2$, along with the input and output to the random oracle, are stored in $\mathcal{H}$. This procedure is also applied in $\mathsf{H}$ if the query number $\ell$ is in $\mathcal{S}$.

Claim 11: For $\boldsymbol{h} \in \mathcal{R}$, let $\mathcal{P} := \mathcal{U}(\mathcal{R}_q)$ and $\mathcal{Q}_{\boldsymbol{h}}$ be the distribution of $\boldsymbol{s}_1 + \boldsymbol{s}_2 * \boldsymbol{h} \mod q$ where $\boldsymbol{s}_1, \boldsymbol{s}_2 \xleftarrow{\$} \mathcal{D}_{\mathcal{R},s}$. Then, for any $a_u \in (1, \infty)$,

$$\Pr[\mathsf{G}_4^{\mathsf{A}} \Rightarrow 1] \leq \left( \max_{(\cdot,\boldsymbol{h}) \in \mathsf{TpdGen}} R_{a_u}(\mathcal{P} \,||\, \mathcal{Q}_{\boldsymbol{h}})^{C_s + L} \cdot \Pr[\mathsf{G}_5^{\mathsf{A}} \Rightarrow 1] \right)^{\frac{a_u - 1}{a_u}}.$$

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

```
$\bar{\mathcal{P}}$
$(\boldsymbol{B}, \boldsymbol{h}) \xleftarrow{\$} \mathsf{TpdGen}$
for $i \in [Q]$
  $\boldsymbol{c}_i \xleftarrow{\$} \mathcal{R}_q$
return $((\boldsymbol{B}, \boldsymbol{h}), \boldsymbol{c}_1, \ldots, \boldsymbol{c}_Q)$
```

```
$\bar{\mathcal{Q}}$
$(\boldsymbol{B}, \boldsymbol{h}) \xleftarrow{\$} \mathsf{TpdGen}$
for $i \in [Q]$
  $(\boldsymbol{s}_1, \boldsymbol{s}_2) \leftarrow \mathcal{D}_{\mathcal{R},s}$
  $\boldsymbol{c}_i := \boldsymbol{s}_1 + \boldsymbol{s}_2 * \boldsymbol{h} \mod q$
return $((\boldsymbol{B}, \boldsymbol{h}), \boldsymbol{c}_1, \ldots, \boldsymbol{c}_Q)$
```

These distributions describe the underlying distributions of $\mathsf{G}_4$ and $\mathsf{G}_5$. By the data processing inequality (Lemma 9) it holds that, for any $a \in (1, \infty)$,

$$R_a(\mathsf{G}_4 \,||\, \mathsf{G}_5) \leq R_a(\bar{\mathcal{P}} \,||\, \bar{\mathcal{Q}}). \tag{3}$$

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

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

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

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

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

By Lemma 10 it follows

$$R_a(\bar{\mathcal{P}} \,||\, \bar{\mathcal{Q}}) \leq \max_{(\cdot,\boldsymbol{h}) \in \mathsf{TpdGen}} R_a(\mathcal{P} \,||\, \mathcal{Q}_h)^Q. \tag{4}$$

Combining probability preservation (Lemma 8) with Equation (3) and Equation (4), we obtain

$$\Pr[\mathsf{G}_5^{\mathsf{A}} \Rightarrow 1] \geq \frac{\Pr[\mathsf{G}_4^{\mathsf{A}} \Rightarrow 1]^{\frac{a}{a-1}}}{R_a(\mathsf{G}_4 \,||\, \mathsf{G}_5)} \geq \frac{\Pr[\mathsf{G}_4^{\mathsf{A}} \Rightarrow 1]^{\frac{a}{a-1}}}{\max_{(\cdot,\boldsymbol{h}) \in \mathsf{TpdGen}} R_a(\mathcal{P} \,||\, \mathcal{Q}_{\boldsymbol{h}})^Q}.$$

The claim follows by setting $Q := C_s + L$ due to at most $C_s$ queries from $\mathtt{Sgn}$ to $\mathsf{H}'$ in Line 32 and $L$ queries to $\mathsf{H}$. ■

*Game* $\mathsf{G}_6$. This game is identical to the previous one except that the output of the preimage sampler $\mathsf{PreSmp}(\boldsymbol{B}, s, (\boldsymbol{c}, \boldsymbol{0}))$ is replaced by a Gaussian over the same lattice, standard deviation and center, namely $\mathcal{D}_{\boldsymbol{\Lambda}(\boldsymbol{B}),s,(\boldsymbol{c},\boldsymbol{0})}$.

Claim 12: For distributions $\mathsf{PreSmp} := \mathsf{PreSmp}(\boldsymbol{B}, s, (\boldsymbol{c}, \boldsymbol{0}))$, $\mathcal{D} := \mathcal{D}_{\boldsymbol{\Lambda}(\boldsymbol{B}),s,(\boldsymbol{c},\boldsymbol{0})}$, and $a_p \in (1, \infty)$ it holds that

$$\Pr[\mathsf{G}_5^{\mathsf{A}} \Rightarrow 1] \leq \max_{(\boldsymbol{B},\cdot)\in\mathsf{TpdGen},\boldsymbol{c}\in\mathcal{R}_q} \left(R_{a_p}(\mathsf{PreSmp} \mid\mid \mathcal{D})^{C_s+1} \cdot \Pr[\mathsf{G}_6^{\mathsf{A}} \Rightarrow 1]\right)^{\frac{a_p-1}{a_p}} .$$

*Proof.* The claim follows by analogous arguments as in Game $\mathsf{G}_5$. Note that we have to replace at most $C_s$ queries in the signing oracle and one additional query after the output of the forgery. ■

*Game* $\mathsf{G}_7$. This game is identical to the previous one except that $\boldsymbol{s}_1, \boldsymbol{s}_2$ are not sampled from a Gaussian distribution centred at $(\boldsymbol{c}, \boldsymbol{0})$ as before. Instead, the preimage of $\boldsymbol{c}$ that was sampled in $\mathsf{H}'$ or $\mathsf{H}$ is used.

Claim 13: $\Pr[\mathsf{G}_6^{\mathsf{A}} \Rightarrow 1] = \Pr[\mathsf{G}_7^{\mathsf{A}} \Rightarrow 1]$.

*Proof.* We need to show that the distributions of the games are the same. The RO output $\boldsymbol{c}$ is the same in both games. In $\mathsf{G}_6$, the signing oracle outputs $(\boldsymbol{s}_1, \boldsymbol{s}_2) \sim \mathcal{D}_{\boldsymbol{\Lambda}(\boldsymbol{B}),s,(\boldsymbol{c},\boldsymbol{0})}$. Since $\boldsymbol{\Lambda}(\boldsymbol{B})$ is the NTRU lattice for $\boldsymbol{h}$ and $q$ and the distribution is shifted by $(\boldsymbol{c}, \boldsymbol{0})$ the output is distributed according to a Gaussian $\mathcal{D}_{\mathcal{R},s}$ conditioned on $\boldsymbol{s}_1 + \boldsymbol{s}_2 * \boldsymbol{h} = \boldsymbol{c} \mod q$. The output distribution in Game $\mathsf{G}_7$ is a Gaussian $\mathcal{D}_{\mathcal{R},s}$ as well where the condition $\boldsymbol{s}_1 + \boldsymbol{s}_2 * \boldsymbol{h} = \boldsymbol{c} \mod q$ is fulfilled by construction (Line 23 and Line 46). ■

*Game* $\mathsf{G}_8$. This game is identical to the previous one except that the game aborts if the $\boldsymbol{s}_1^\star, \boldsymbol{s}_2^\star$ corresponding to the adversary's output equals the preimage with which the challenge random oracle output $\boldsymbol{c}^\star$ was computed and the message/signature pair is not in $\mathcal{Q}$.

Claim 14: $\left|\Pr\left[\mathsf{G}_7^{\mathsf{A}} \Rightarrow 1\right] - \Pr\left[\mathsf{G}_8^{\mathsf{A}} \Rightarrow 1\right]\right| \leq \frac{1}{2^{n-1}}$.

*Proof.* We distinguish two cases: first, the adversary queried the signing oracle corresponding to $\boldsymbol{c}^\star$, i.e. a signing query to $m^\star$ output $(r^\star, \cdot)$ and $\mathsf{H}(\boldsymbol{h}, r^\star, m^\star) = \boldsymbol{c}^\star$. In this case, the abort event cannot trigger because if the preimages are the same, the signature must be the same as well and therefore $(m^\star, \sigma^\star) \in \mathcal{Q}$. Second, if the signing oracle was not queried corresponding to $\boldsymbol{c}^\star$, the adversary does not have any information about the preimages of $\boldsymbol{c}^\star$ except that they are Gaussian distributed. Hence, we obtain the claimed upper bound by using the min-entropy of a sample from a Gaussian distribution conditioned on $\boldsymbol{c}^\star$ from Lemma 6. ■

*Reduction from $\mathcal{R}$-**SIS** to* $\mathsf{G}_8$. We now can reduce $\mathcal{R}$-**SIS** to Game $\mathsf{G}_8$.

Claim 15: There exists an adversary $\mathsf{B}$ against $\mathcal{R}$-**SIS** such that

$$\Pr[\mathsf{G}_8^{\mathsf{A}} \Rightarrow 1] \leq \mathrm{Adv}_{1,q,\alpha,2\beta,\mathsf{B}}^{\mathcal{R}\text{-}\mathbf{SIS}}.$$

*Proof.* Adversary $\mathsf{B}$ is formally constructed in Figure 8. Due to the changes in the previous games, adversary $\mathsf{B}$ can perfectly simulate the game for adversary $\mathsf{A}$ against $\mathsf{G}_8$ without having the secret key for $\boldsymbol{h}$. Let us assume, that $\mathsf{A}$ wins the strong unforgeability game $\mathsf{G}_8$, i.e. the forgery verifies and the tuple $(m^\star, \sigma^\star)$ was not queried before. That means that the output of adversary $\mathsf{B}$ fulfills the following conditions. First, it holds that

$$(\boldsymbol{s}_1^\star - \boldsymbol{s}_1') + (\boldsymbol{s}_2^\star - \boldsymbol{s}_2') * \boldsymbol{h} = \boldsymbol{s}_1^\star + \boldsymbol{s}_2^\star * \boldsymbol{h} - (\boldsymbol{s}_1' + \boldsymbol{s}_2' * \boldsymbol{h}) = \boldsymbol{c}^\star - \boldsymbol{c}^\star = \boldsymbol{0}$$

due to the computation of $s_1^\star$ in Line 13 and the structure of elements in $\mathcal{H}$. Second, it cannot equal $\boldsymbol{0}$ due to the changes in $\mathsf{G}_8$. Third, the norm bound of the output can be upper bounded by $B = 2\beta$:

$$\|(\boldsymbol{s}_1^\star - \boldsymbol{s}_1', \boldsymbol{s}_2^\star - \boldsymbol{s}_2')\|_2 \leq \|(\boldsymbol{s}_1^\star, \boldsymbol{s}_2^\star)\|_2 + \|(\boldsymbol{s}_1', \boldsymbol{s}_2')\|_2 \leq 2\beta,$$

where the last inequality follows by the winning condition of adversary $\mathsf{A}$ and the norm condition of preimages since the game does not abort in Line 09.

■

■

```
$\mathsf{B}(\boldsymbol{h})$
$\mathcal{H}, \mathcal{Q} \leftarrow \emptyset$
$cnt, \ell := 0$
choose random $\mathcal{S} \subseteq [Q_{\mathsf{H}} + 1]$ such that $|\mathcal{S}| = L$
$(m^\star, \sigma^\star) \xleftarrow{\$} \mathsf{A}^{\mathtt{Sgn}(\cdot),\mathsf{H}(\cdot,\cdot,\cdot)}(\boldsymbol{h})$
**parse** $\sigma^\star \rightarrow (r^\star, \boldsymbol{s}_2^\star)$
$\boldsymbol{c}^\star := \mathsf{H}(\boldsymbol{h}, r^\star, m^\star)$
**find** $(\boldsymbol{s}_1', \boldsymbol{s}_2') : (\boldsymbol{c}^\star, \boldsymbol{h}, r^\star, m^\star, \cdot, \boldsymbol{s}_1', \boldsymbol{s}_2') \in \mathcal{H}$
**if** $\|(\boldsymbol{s}_1', \boldsymbol{s}_2')\|_2 > \beta$
  **abort**
**find** $\ell^\star : (\boldsymbol{c}^\star, \boldsymbol{h}, r^\star, m^\star, \ell^\star, \cdot, \cdot) \in \mathcal{H}$
**if** $\ell^\star \notin \mathcal{S} \wedge \ell^\star \neq \perp$
  **abort**
$\boldsymbol{s}_1^\star := \boldsymbol{c}^\star - \boldsymbol{s}_2^\star * \boldsymbol{h} \mod q$
**if** $(\boldsymbol{s}_1^\star, \boldsymbol{s}_2^\star) = (\boldsymbol{s}_1', \boldsymbol{s}_2') \wedge (m^\star, \sigma^\star) \notin \mathcal{Q}$
  **abort**
**return** $(\boldsymbol{s}_1^\star - \boldsymbol{s}_1', \boldsymbol{s}_2^\star - \boldsymbol{s}_2')$

**Oracle** $\mathsf{H}(pk, r, m)$
**return** $\mathsf{G}_8.\mathsf{H}(pk, r, m)$

**Oracle** $\mathsf{H}'(\boldsymbol{h}, r, m)$
**return** $\mathsf{G}_8.\mathsf{H}'(\boldsymbol{h}, r, m)$

**Oracle** $\mathtt{Sgn}(m)$
**return** $\mathsf{G}_8.\mathtt{Sgn}(m)$
```

**Figure 8.** Adversary $\mathsf{B}$ against $\mathcal{R}$-**SIS** for the proof of Theorem 1.

# C Appendix for Section 6

```
sage: SIS.estimate.rough(SIS.Parameters(n=512,q=12289,length_bound=5833.93,norm=2,m=2*512))
lattice :: rop: ≈2^121.2, red: ≈2^121.2, δ: 1.003882, β: 415, d: 1024, tag: euclidean
sage: SIS.estimate.rough(SIS.Parameters(n=512,q=12289,length_bound=2*5833.93,norm=2,m=2*512))
lattice :: rop: ≈2^95.8, red: ≈2^95.8, δ: 1.004561, β: 328, 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: ≈2^279.2, red: ≈2^279.2, δ: 1.002114, β: 956, d: 2048, tag: euclidean
```

**Figure 9.** SIS hardness estimates for ring dimension $n = 512$, $n = 1024$ and length bound $\beta$, $2\beta$.

## C.1 Proof of Lemma 17

**Lemma 17 (Optimal $C_s$).** For Falcon$^+$-512 with $\lambda = 128$ it holds that,

$$\operatorname*{arg\,min}_{C_s} \left\{ C_s \,\middle|\, \sum_{i=0}^{Q_s} \binom{C_s}{i} (1 - p_{\mathsf{PreSmp},\beta})^{C_s - i} (p_{\mathsf{PreSmp},\beta})^i \leq 2^\lambda \right\} \leq 2^{64} + 2^{50},$$

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

$$\operatorname*{arg\,min}_{C_s} \left\{ C_s \,\middle|\, \sum_{i=0}^{Q_s} \binom{C_s}{i} (1 - p_{\mathsf{PreSmp},\beta})^{C_s - i} (p_{\mathsf{PreSmp},\beta})^i \leq 2^\lambda \right\} \leq 2^{64} + 2^{36}.$$

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

$$\begin{aligned}
p_{\mathsf{PreSmp},\beta} &= \max_{\substack{\boldsymbol{c} \in \mathcal{R}_q \\ (\boldsymbol{B},\cdot) \in \mathsf{sup}(\mathsf{Gen})}} \Pr_{(\boldsymbol{s}_1,\boldsymbol{s}_2) \xleftarrow{\$} \mathsf{PreSmp}(\boldsymbol{B},s,(\boldsymbol{c},\boldsymbol{0}))} [\|(\boldsymbol{s}_1,\boldsymbol{s}_2)\|_2 \leq \beta] \\
&= \max_{\substack{\boldsymbol{c} \in \mathcal{R}_q \\ (\boldsymbol{B},\cdot) \in \mathsf{sup}(\mathsf{Gen})}} 1 - \Pr_{(\boldsymbol{s}_1,\boldsymbol{s}_2) \xleftarrow{\$} \mathsf{PreSmp}(\boldsymbol{B},s,(\boldsymbol{c},\boldsymbol{0}))} [\|(\boldsymbol{s}_1,\boldsymbol{s}_2)\|_2 > \beta] \\
&\geq \max_{\substack{\boldsymbol{c} \in \mathcal{R}_q \\ (\boldsymbol{B},\cdot) \in \mathsf{sup}(\mathsf{Gen})}} 1 - \Pr_{(\boldsymbol{s}_1,\boldsymbol{s}_2) \leftarrow \mathcal{D}_{\boldsymbol{\Lambda}(\boldsymbol{B}),s,(\boldsymbol{c},\boldsymbol{0})}} [\|(\boldsymbol{s}_1,\boldsymbol{s}_2)\|_2 > \beta] \cdot (1 + 2\epsilon) && \text{(Lemma 12, Lemma 15)} \\
&\geq \max_{\substack{\boldsymbol{c} \in \mathcal{R}_q \\ (\boldsymbol{B},\cdot) \in \mathsf{sup}(\mathsf{Gen})}} 1 - \left( \frac{\rho_s(\boldsymbol{\Lambda}(\boldsymbol{B}))}{\rho_{s,(\boldsymbol{c},\boldsymbol{0})}(\boldsymbol{\Lambda}(\boldsymbol{B}))} \cdot \left(\sqrt{e^{1-\tau^2}\tau^2}\right)^{2n} \cdot (1 + 2\epsilon) \right) && \left(\text{Lemma 2 and } \beta = \frac{\tau s \sqrt{2n}}{\sqrt{2\pi}}\right) \\
&\geq 1 - \left( \frac{1+\epsilon}{1-\epsilon} \cdot \left(\sqrt{e^{1-\tau^2}\tau^2}\right)^{2n} \cdot (1 + 2\epsilon) \right). && (\text{Lemma 5 and } s \geq \eta_\epsilon(\boldsymbol{\Lambda}(\boldsymbol{B})))
\end{aligned}$$

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

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

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

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

When the following condition is satisfied:

$$Q_s \leq C_s p_{\mathsf{PreSmp},\beta}, \tag{5}$$

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,

$$\sum_{i=0}^{Q_s} \binom{C_s}{i} (1 - p_{\mathsf{PreSmp},\beta})^{C_s - i} (p_{\mathsf{PreSmp},\beta})^i \leq \exp\left(-2C_s \left(p_{\mathsf{PreSmp},\beta} - \frac{Q_s}{C_s}\right)^2\right) \tag{6}$$

where $Q_s$ is the number of successes, $C_s$ is the number of trials, and $p_{\mathsf{PreSmp},\beta}$ is the probability of success in each trial. To satisfy the condition of Equation (5), $C_s$ is set as follows,

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

Finally, the bound in Equation (6) is verified as follows,

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

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

■

### C.2 Proof of Corollary 3

**Corollary 3 (Rényi Loss for Falcon$^+$-512 (Preimage Sampler) in Thm. 2).** For $\varepsilon \geq 2^{-\lambda} = 2^{-120}$, $r_p = R_{a_p}(\mathsf{PreSmp} \parallel \mathcal{D})$, $C_s = 2^{64} + 2^{50}$, 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 $a_p \approx 72.96$.

*Proof.* By Lemma 16 we need to solve

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

By Corollary 1 we can upper bound $R_{a_p}$

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

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

$$a_p = \frac{\lambda\, \epsilon^2 \ln(4) + \sqrt{8C_s\, \lambda\, \epsilon^2 \ln(2) + \lambda^2\, \epsilon^4 \ln^2(4)}}{4C_s\, \epsilon^2} \tag{7}$$

Plugging $\lambda = 120$, $\epsilon = 1/\sqrt{2^{64} \cdot 128} = 2^{-35.5}$ and $C_s = 2^{64} + 2^{50}$ into Equation (7) gives

$$a_p \approx 72.96$$

and thus a bit loss of at most

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

■

### C.3 Proof of Corollary 4

**Corollary 4 (Rényi Loss for Falcon$^+$-512 (Uniformity) in Thm. 2).** For $\varepsilon \geq 2^{-\lambda} = 2^{-116.5}$, $r_u = R_{a_u}(\mathcal{U}(\mathcal{R}_q) \,||\, \mathcal{U}_{\boldsymbol{h}})$, $C_s = 2^{64} + 2^{50}$, 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.* Similar to the proof of Corollary 3 except that the Rényi divergence is upper bounded using Corollary 2. So we get

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

■