Quantum Oracle Distribution Switching and Applications to Falcon and Ring Signatures

Marvin Beckmann[1] and Christian Majenz[1]

Technical University of Denmark, Kongens Lyngby 2800 {mabeck, chmaj}@dtu.dk

Abstract. Motivated by digital signature algorithms ranging from Falcon to fully-anonymous ring signatures used in Signal-style key exchange, such as Gandalf, we revisit a fundamental problem in post-quantum security proofs: distinguishing oracle functions whose outputs are sampled independently from distributions P and Q that are close. In the context of signatures, closeness is often measured via Rényi divergences, which yield *multiplicative* bounds in the classical setting. A counterexample shows that such multiplicative-error bounds for distinguishers with classical oracle access have no analogue for quantum access, and we provide two alternative approaches based on small-range distributions and reprogramming techniques. We also give a concrete, optimal bound for the case where P and Q are close in statistical distance.
We apply these techniques to the motivating constructions. (i) We give the first QROM security proof for Falcon that avoids oracle indistinguishability arguments based on statistical distance. This is crucial, as Falcon's ROM proof relies on Rényi divergence, while the statistical distance induced by its parameters is too large to yield meaningful bounds. (ii) We formalize and abstract the ring signature construction used in Gandalf as a modular framework by defining ring trapdoor preimage-sampleable functions (RPSFs), for which we obtain two QROM proofs. We also provide two QROM security proofs for AOS ring signatures, adapting existing QROM techniques. Together with our results on RPSF-based ring signatures, this yields QROM security proofs for a broad class of fully-anonymous linear ring signature constructions, including Gandalf and the AOS-based constructions Erebor and MayoRS.

Keywords: QROM · Rényi divergence · Falcon · Ring Signatures · Oracle Indistinguishability

Table of Contents

Quantum Oracle Distribution Switching and Applications to Falcon and Ring Signatures . . . 1
Marvin Beckmann and Christian Majenz
1 Introduction . . . 4
2 Preliminaries . . . 9
2.1 Signatures . . . 10
2.2 Sigma Protocols . . . 10
2.3 QROM . . . 10
2.4 Deutsch-Jozsa Algorithm . . . 13
3 Quantum Oracle Distribution Switching . . . 13
3.1 Oracle Switching using Statistical Distance . . . 13
3.2 Oracle Switching using Rényi Divergence . . . 14
3.3 Reprogramming with Different Distributions . . . 16
4 A Closer Look at Falcon in the QROM . . . 17
5 Ring Signature from Ring Preimage Sampleable Functions . . . 20
5.1 Unforgeability via Adaptive Reprogramming . . . 21
5.2 Unforgeability via History-Free Proofs . . . 24
5.3 Anonymity . . . 24
6 AOS Ring Signatures . . . 25
6.1 SUF-CRA to UF-NRA . . . 26
6.2 UF-NRA: General Strategy using Measure-and-Reprogram . . . 27
6.3 UF-NRA: Commit-and-Open Sigma Protocols . . . 28
7 Instantiation and Discussion of Ring Signatures . . . 32
A Additional Preliminaries . . . 39
A.1 Probabilities . . . 39
A.2 Signatures . . . 39
A.3 Ring Signatures . . . 40
A.4 Sigma Protocols . . . 43
A.5 Quantum Computation Basics . . . 45
A.6 NTRU Lattices . . . 45
B Omitted Proofs . . . 48
B.1 Proof of Theorem 3 . . . 48
B.2 Proof of Proposition 1 . . . 53
B.3 Proof of Lemma 6 . . . 55
B.4 Proof of Theorem 8 . . . 56
B.5 Proof of Theorem 9 . . . 59
B.6 Proof of Lemma 8 . . . 60
B.7 Proof of Lemma 10 . . . 60
C History-Freeness for Ring Signatures . . . 62
C.1 Proof of Theorem 7 . . . 66
D AOS Ring Signatures from Merkle Tree Based C&O Protocols . . . 67

E Applications and Instantiations 68
E.1 Falcon in the QROM 69
E.2 Application to Gandalf 73

1 Introduction

A digital signature scheme allows a user to authenticate messages to other parties. With the growing threat of quantum computers, post-quantum (PQ) secure signatures are essential. National Institute of Standards and Technology (NIST) concluded its PQ standardization process with three signature standards: ML-DSA, FN-DSA, and SLH-DSA. Of these, FN-DSA (formerly Falcon [33]) has received significant attention in recent cryptographic research.

At its core, Falcon is built on the GPV framework [36], a general paradigm for constructing hash-and-sign signatures from preimage sampleable functions. However, Falcon deviates from the GPV framework in two important ways. First, it uses rejection sampling (aborts and repeats) to ensure correctness. Second, and more relevant to this work, it employs optimization techniques that replace statistical distance arguments with Rényi divergence properties, enabling significantly tighter parameter choices. As shown in [34], the statistical distance of Falcon's preimage sampler is only on the order of 2^{-34}, which is far too large for standard proof techniques.

The existing quantum random oracle model (QROM) proof for the GPV framework [12] proceeds by simulating the entire random oracle with hash values that are statistically close to uniform. Giving an adversary quantum access to such a simulated oracle remains indistinguishable from the honest oracle, up to an error depending on the statistical distance and the number of queries. However, this approach fundamentally requires the simulated hash values to have negligible statistical distance from uniform. Since Falcon's optimization techniques replace statistical distance with Rényi divergence, these existing QROM techniques do not apply, and there is currently no argument that Falcon with its current parameters is secure in the QROM.

Closing this gap requires understanding how quantum oracle access behaves when the underlying output distribution is switched from one distribution to another that is close in terms of Rényi divergence rather than statistical distance. Classically, the Rényi divergence yields a clean purely multiplicative bound [47,30] when switching from the uniform to the simulated distribution. Whether an analogous bound holds if provided quantum oracle access is the central technical question we address in this work.

Rivest, Shamir, and Tauman [60] introduced ring signature schemes (RSSs), which provide the same functionality as group signatures without trusted setup, allowing a member to sign on behalf of a group without revealing their identity. The same GPV framework that underlies Falcon also serves as the foundation for ring signature constructions like Gandalf [35] and FalconRS, which extend the preimage sampleable function approach to the ring setting.

A prominent application of RSSs is the construction of deniable authenticated key exchange (DAKE), e.g., for messenger applications. The most widely used key exchange protocol for end-to-end-encrypted instant messaging is the Signal protocol (WhatsApp, Signal, and Facebook Messenger, etc.). Its initial DAKE, X3DH [52], is, however, based on the Diffie-Hellman key exchange, which can be broken by quantum computing attacks. A variant of that protocol, PQXDH [46],

uses a hybrid approach with ML-KEM to ensure PQ confidentiality. PQXDH thus prevents "harvest-now-decrypt-later" attacks, but lacks PQ authentication and might not be PQ-anonymous either ("harvest-now-judge-later" [45]). Most proposals for fully PQ secure Signal-conforming DAKEs use RSS [41,15,42]. An authenticated KEM (AKEM) is also the primitive behind two modes of the hybrid public key encryption standard [6]. Recent work [35] provides a generic construction of deniable AKEMs based on KEMs and RSSs.

Signal-conforming DAKE constructions apply RSSs for ring size 2, and require strong anonymity (anonymity under full key-exposure) [8]. In this context, RSSs with a signature size growing linearly in the ring size are typically more efficient than logarithmic-sized RSSs. Besides constructions like Gandalf, there is a second approach that uses the AOS-transform [1], transforming Σ-protocols into linear-sized RSSs, used in Erebor and MayoRS [13,45].

Both approaches and the resulting explicit constructions are currently only proven secure in the classical random oracle model (ROM), and thus *not* supported by a QROM proof. In other words, they do not enjoy provable PQ security. Prior to our work, there were thus two options for PQ Signal-conforming DAKE: Accepting the lower level of assurance provided by a ring-signature-based protocol without provable PQ security, or using a split-KEM-based protocol and accepting its disadvantages (like its inefficiency).

Our Contributions. In this work, we provide QROM proofs for Falcon and the two types of RSS used in Signal-conforming DAKEs.

To facilitate the QROM proofs that use Rényi divergence-based optimizations, we give a comprehensive set of results for using Rényi divergences for the (in)distinguishability of quantum-accessible oracles. This includes a counterexample showing that the main approach used in the classical setting cannot work for quantum-accessible oracles, as well as two alternatives that do work: Targeted reprogramming of output values for high-entropy input based on adaptive reprogramming results [38], and a technique to replace all outputs via small-range distributions [64].

For the ring trapdoor paradigm of constructing RSSs, we generalize and abstract the approach of FalconRS and Gandalf by formalizing a novel primitive: ring preimage sampleable functions (RPSFs). We then give a generic RSS construction from RPSFs that admits security proofs using either statistical distance or Rényi divergence properties independently for domain sampling and preimage sampling, yielding four concrete security bounds. Finally, we prove QROM security for the generic RPSF-based RSS construction using two different techniques. One of them, (a generalization of) the formalism of history-free reductions [12], is relatively straightforward and relies on statistical distance arguments. The other uses Rényi divergences for tighter parameters, and requires use of our study of Rényi divergences for quantum-accessible oracles.

For AOS RSSs, we give two security bounds, one for generic Σ-protocols and a tighter bound for Σ-protocols with commit-and-open structure. To obtain the generic bound, we employ the measure-and-reprogram technique [26]. While the resulting bound grows exponentially with the ring size, it can give meaningful

guarantees for ring size 2, as needed for the application to Signal-conforming DAKEs. For the second bound, practical for larger rings, we use straightline extraction techniques, yielding a multiplicatively tight reduction to the special soundness of the Σ-protocol. The result can be extended to Merkle tree-based commit-and-open Σ-protocols.

Technical Overview. In the following, we give a more detailed overview.
Quantum Oracle Distribution Switching. Consider an adversary with quantum access to an oracle for a function $f_P : \mathcal{X} \to \mathcal{Y}$. The outputs of f are independently sampled from a distribution P over $\mathcal{Y}$ for each $x \in \mathcal{X}$. If P is the uniform distribution, then this is exactly the QROM. We give several results on the distinguishability of such oracles for pairs of distributions P and Q.

First, we provide an asymptotically tight *explicit* bound on the statistical distance of an algorithm's output, where the algorithm interacts with an oracle for f_P or f_Q, for two distributions P and Q, as a function of the statistical distance of P and Q.[1] For this explicit bound, we use a *compressed oracle* [66] view, and analyze the trace norm of the final states produced by an algorithm interacting with the two oracles. By adding null-terms, we bound the trace norm by a sum of the operator norms of two "compression" operators that depend on the underlying classical distributions. The norm of these operators can be bounded by the statistical distance of the classical distributions.

Now consider the Rényi divergence $R_\alpha(P\|Q)$ of order α between probability distributions P and Q. For an algorithm making q classical queries to f_P or f_Q, the probability of any output-dependent event under P admits a multiplicative bound in terms of the Rényi divergence as $R_\alpha(P\|Q)^q$ with the event occuring under Q, up to an additional exponent $\frac{\alpha-1}{\alpha}$.

We show that no analogous bound can hold in the quantum setting. Our construction proceeds as follows. From f_R, where $R \in \{P, Q\}$, we define a Boolean function $h[f_R]$ on a domain of size 2^{1+n} such that: (i) a single query to $h[f_R]$ can be implemented using 2 queries to f_R; and (ii) $h[f_Q]$ is *balanced in expectation* over the choice of f_Q, whereas $h[f_P]$ is biased in expectation.

We then apply the Deutsch-Jozsa algorithm [24] to $h[f_R]$ using a single query. This yields a distinguisher which, given oracle access to f_R, outputs 1 whenever the algorithm reports "constant". Since $h[f_Q]$ is balanced in expectation, this occurs with probability $\mathcal{O}(2^{-n})$ for $R = Q$, while it occurs with constant probability for $R = P$. Hence, their ratio grows as $\Theta(2^n)$, which outpaces any loss depending only on P and Q, and continues to do so even after weakening the comparison to $\Pr\left[1 \leftarrow \mathcal{A}^{f_Q}\right]^\gamma$ for any fixed $\gamma > 0$.

As a positive result, we show that by accepting an additive error, we can use the *small-range distribution* toolkit from [64]. Here, r values are sampled according to the underlying distribution, and for each value in $\mathcal{X}$, one of these r samples is assigned as the value in $\mathcal{Y}$. For these r samples, we can use the

[1] Asymptotically, this was also shown in [7], but without explicit bounds which are needed for parameter determination in practice. We need explicit bounds, so we included the theorem statement with explicit bounds in the main body, and the proof in Appendix B.1.

classical Rényi divergence properties and bound the overall capabilities of any algorithm. However, the error term depends on r, and for practical applications, r must be chosen super-polynomially large, increasing the power of the Rényi divergence factor. To have practical application, it turns out that the statistical distance between the two distributions would also need to be negligible, so the Rényi divergence would not need to be used in the first place. This concludes the Rényi divergence study with the insight that replacing the entire underlying oracle is impractical when using the Rényi divergence.

Taking a closer look at explicit constructions, we observed that the hash input usually includes a salt to achieve strong unforgeability. In this case, we can use results on *adaptive reprogramming* [38], also known as *resampling* [2], to *replace only the outputs used by the signing oracle*. In this technique, R positions are reprogrammed by sampling them again from the same underlying distribution before they are programmed into the oracle. Replacing these samples with samples from a different distribution, we can use the classical properties of the Rényi divergence (or the statistical distance).

Falcon. Falcon has enough entropy in the message hash input to enable our results for adaptive reprogramming in combination with the Rényi divergence. The reduction follows ideas from [34] closely. First, we reprogram the hashes from each preimage sampling try using one of the C_s targets from the C_s-$\mathcal{R}$-NTRU-SPISIS* assumption, where C_s bounds the total number of preimage sampling attempts across all signing queries. This assumption provides C_s syndromes together with Gaussian preimages and asks the adversary to find a *new* short preimage for any of them. We apply our results on the Rényi divergence here. In the next step these signatures are replaced by the preimages provided with the hash values in the previous game. Finally, we can either include challenges from a multi-target assumption t-$\mathcal{R}$-NTRU-ISIS using small-range distributions or apply measure-and-reprogram to include a challenge from a single-target assumption $\mathcal{R}$-NTRU-ISIS. The assumptions differ in that they have either t or 1 uniform targets for which a valid preimage needs to be found.

Ring Trapdoor Ring Signatures. We introduce ring trapdoor function (RTDF) and describe preimage sampleable properties for them, generalizing preimage sampleable (trapdoor) functions (PSFs) [36], to give a general RSS construction. Brakerski and Kalal [14] introduce a less general framework from PSFs to construct ring signatures in the standard model. The first digital signature scheme for which QROM security was proven uses PSFs [12]. Unsurprisingly, similar strategies can be employed when proving the security of our RSS: We observe that a *history-free* ROM proof can be given. For this, we adapt the original notion of history-free proofs for digital signatures to the setting of RSSs. In essence, we can simulate the random oracle (RO) by composing a *private* RO with the domain sampler and the evaluation function of the RTDF. Combining the domain sampler and the evaluation function to simulate the RO requires statistical distance arguments, and under those, this mapping produces a distribution in the range close to uniform, replacing the RO. Similar to the ordinary signature case, the classical history-free reduction implies a QROM reduction.

We present a second separate result that works with properties based on the Rényi divergence that uses similar strategies as the proof for FALCON.

AOS Ring Signatures. Here, we first describe how to simulate the signing oracle. This again requires adaptive reprogramming in the QROM. To simulate, replace the single instance of the honest Σ-protocol prover by the honest verifier zero knowledge (HVZK) simulator, and reprogram the RO accordingly. Given sufficient min-entropy in the simulated commitments and thus the RO inputs, we can apply the adaptive reprogramming lemma from [38]. It is from this point on that we give two reductions, one for generic sigma protocols and a tight one for commit-and-open protocols.

The strategy for the generic proof is to construct an impersonation adversary against the underlying Σ-protocol using a successful forgery adversary. This adversary must send a commitment com, receive a challenge ch, and send a response rsp in order. Like in ordinary Fiat-Shamir, this requires programming the random oracle, but here, challenges are computed from the previous commitments in the ring. In the impersonation game, the challenge is received after the commitment, so to reprogram, the query producing the challenge has to happen *after* the commitment was queried, requiring a definite time order. To get such an order, we use the measure-and-reprogram technique [26]. This technique incurs a multiplicative loss in $(2q+1)^{2n}$, where n is the number of queries for which the order is desired. As we need an order of all queries for the ring, and the ring is of size N, this introduces an exponential loss in N. Here, one of the reprogrammings actually programs the Σ-protocol challenge, and the remaining ones *commit* the adversary to a *classical ordering* of the queries made for producing the signature.

For commit-and-open protocols, we use a different approach. Commit-and-open protocols in the (Q)ROM allow RO-based extraction of commitments. If an adversary succeeds for sufficiently many different challenges, one can use (generalized) special soundness to extract a witness. Generalizing the strategy from [19,28], we define a compressed oracle database property for fooling the extractor, i.e., the database contains a valid forgery, but the special soundness extractor fails. Using the quantum transition capacity framework [20] we bound the probability that after q_{H} hash queries, this database property holds. If the database property does not hold, and the prover succeeds, we can measure the compressed oracle database to recover a witness. An additional Merkle tree optimization only requires slight modifications in the database properties.

Additional Related Work. To the best of our knowledge, the AOS framework has not been analyzed in the QROM, and ring trapdoor constructions also have not been proven secure in the QROM. There is one work [23] that considers QROM security of a logarithmic-size ring signature based on symmetric primitives to build an accumulator and simulation-sound extractability of an additional proof system (not considering inefficient plain model RSSs like [14]). The specific ring signatures, though logarithmic in asymptotics, are in the order of several hundred KBs even for small rings, and therefore practically infeasible. Recent candidates for fully anonymous RSSs [51,11,13,49,63] are also not proven

secure in the QROM. We identified some works that consider the QROM, but they are either not fully anonymous [40], or consider QROM quantum-access-secure RSS [17,18], but are inefficient. The framework of [21] for group signatures and accountable RSSs covers neither the AOS transform nor the ring-trapdoor paradigm, but discarding the msk in accountable RSSs does yield anonymity under full key exposure. This does rely on assuming the msk is irrecoverable, and the repetition parameter ($2\lambda N$) in their work puts signature sizes out of practical range for Signal-conforming DAKEs. Katsumata et al. [45] present a new weaker anonymity notion for RSSs that they argue to be sufficient. Full anonymity implies this new deniability notion. Signal-conforming DAKEs can also be constructed from split-KEMs like [22,56] instead of RSS using the K-WAAY [22] protocol.

Concurrent independent work. After the research for this paper was concluded, we noticed a concurrent independent work [54] that also generalizes the RTDF approach from [14] to capture the work of GANDALF. Their work does not consider QROM security, and they consider weaker anonymity notions.

Subsequent to our analysis of the AOS RSS, there has been some further work on ring signatures, specifically an efficient QROM secure ring signature PEGARING [67] was introduced whose security is also proven using commit-and-open Σ-protocols in the DualRing [63] framework using techniques similar to those we applied for proving the security of the AOS framework.

2 Preliminaries

By $a \leftarrow \mathcal{D}$, we assign a value to a given $\mathcal{D}$, where $\mathcal{D}$ is a distribution or an algorithm. If D is a finite set, $a \leftarrow D$ denotes uniform sampling from the set. Let $\mathcal{U}(D)$ denote the uniform distribution over D. We denote $[n] := \{1, 2, \ldots, n\}$. For standard definitions of the statistical distance $\Delta(P, Q)$, the Rényi divergence $R_\alpha(P\|Q)$ and the Kullback-Leibler divergence $D_{KL}(P\|Q)$, see Appendix A.1.

Lemma 1 ([47, Lemma 4.1] and [30, Theorem 1 and 9]). *Let P and Q be two discrete probability distributions and E an event such that $E \subseteq \mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q)$. Let $f : \mathrm{Supp}(Q) \to \mathcal{X}$ be a function (or stochastic map). For any $\alpha \in (1, \infty]$, we have the probability preservation property and the data processing inequality,*

$$\Pr[P \in E]^{\frac{\alpha}{\alpha-1}} \leq R_\alpha(P\|Q) \cdot \Pr[Q \in E], \quad \text{and} \quad R_\alpha(f(P)\|f(Q)) \leq R_\alpha(P\|Q).$$

Lemma 2 ([3, Lemma 2.9]). *Let $\alpha \in (1, \infty]$, and let $P_1, \ldots, P_q$ and $Q_1, \ldots, Q_q$ be distributions with $\mathrm{Supp}(P_i) \subseteq \mathrm{Supp}(Q_i)$. For independent coordinates it holds that*

$$R_\alpha\left(\bigotimes_{i\in[q]} P_i \middle\| \bigotimes_{i\in[q]} Q_i\right) = \prod_{i\in[q]} R_\alpha(P_i\|Q_i) \leq \left(\max_{i\in[q]} R_\alpha(P_i\|Q_i)\right)^q.$$

Lemma 3 ([58, Lemma 1]). *Let $\mathcal{A}^P, \mathcal{A}^Q$ be algorithms making at most q queries to an oracle sampling from distribution P and Q respectively, and returning a bit. Let $0 \leq D_{KL}(P\|Q) \leq \epsilon$. Then, it holds that*

$$\left|\Pr\left[1 \leftarrow \mathcal{A}^P\right] - \Pr\left[1 \leftarrow \mathcal{A}^Q\right]\right| \leq \sqrt{q\epsilon/2}.$$

2.1 Signatures

The standard definitions of digital signatures are recalled in Section A.2.

A Ring signature scheme (RSS) is a quadruple of probabilistic polynomial time (ppt) algorithms $(\mathsf{Stp}, \mathsf{KGen}, \mathsf{Sign}, \mathsf{Vf})$. The key generation is performed by each participant individually. A ring ρ consists of the individual public keys of at most κ participants. For this ring, each member of ρ can generate a signature σ on behalf of the entire ring using their own secret key. A signature can be verified using only the public information ρ. The formal syntax and security definitions can be found in Appendix A.3.

2.2 Sigma Protocols

A Σ-protocol is a 3-round public-coin interactive proof $(\mathsf{P} = (\mathsf{P}_1, \mathsf{P}_2), \mathsf{V} = (\mathsf{V}_1, \mathsf{V}_2))$ for a relation $R \subseteq \mathcal{I} \times \mathcal{W}$: First, the prover P_1 sends a *commitment* com; then the verifier V_1 responds with a random *challenge* $\mathsf{ch} \in \mathcal{C}$; and finally, P_2 sends a *response* rsp. This final response is evaluated using V_2. The formal definition and well-known properties can be found in Appendix A.4.

2.3 QROM

Our main technical proofs rely on reprogramming techniques [26,38] and on a framework [20] for proving query complexity bounds in the QROM. The latter is a framework exploiting Zhandry's compressed-oracle technique. We use a slightly adjusted version of the framework by [28]. We also use a proof technique for digital signatures from [12] based on history-free reductions, which we generalize to the ring signature setting. We model quantum access to a random oracle $\mathsf{O} : \mathcal{X} \to \mathcal{Y}$ via oracle access to a unitary $\boldsymbol{U}_\mathsf{O}$ defined by $|x\rangle_\mathcal{X}|y\rangle_\mathcal{Y} \mapsto |x\rangle_\mathcal{X}|y \oplus \mathsf{O}(x)\rangle_\mathcal{Y}$, and adversaries with quantum access to O act as a sequence of unitaries, interleaved with applications of $\boldsymbol{U}_\mathsf{O}$. We utilize two reprogramming techniques.

Measure-and-reprogram. The first, *measure-and-reprogram* [27,26], allows measuring randomly selected query inputs and reprogramming the random oracle at those inputs to fresh random outputs. If an algorithm outputs a tuple of query inputs such that the resulting input-output pairs together with an additional output fulfill a predicate, then the same will hold when the measurements and reprogrammings are applied, with respect to the reprogrammed values. This also yields the order in which the relevant queries were made.

Theorem 1 ([26, Theorem 6]). *Let n be a positive integer, and let $\mathcal{X}$ and $\mathcal{Y}$ be finite non-empty sets. Let $\mathcal{A}$ be an arbitrary oracle quantum algorithm that*

makes q queries to a uniformly random $\mathsf{H} : \mathcal{X} \to \mathcal{Y}$ and that outputs a tuple $\boldsymbol{x} = (x_1, \ldots, x_n) \in \mathcal{X}^n$ and a (possibly quantum) output z. There exists a black-box polynomial-time $(n+1)$-stage quantum algorithm $\mathcal{S}$, satisfying the following properties. $\mathcal{S}$ has the following syntactic behavior: in the first stage it outputs a permutation π together with $x_{\pi(1)}$ and takes as input $\Theta_{\pi(1)}$, and then for every subsequent stage $1 < i \leq n$ it outputs $x_{\pi(i)}$ and takes as input $\Theta_{\pi(i)}$; eventually, in the final stage (labeled by $n+1$) it outputs z. We denote such an execution of $\mathcal{S}$ as $(\pi, \pi(\mathbf{x}), z) \leftarrow \langle \mathcal{S}^{\mathcal{A}}, \pi(\boldsymbol{\Theta}) \rangle$. For any $\boldsymbol{x}' \in \mathcal{X}^n$ without duplicate entries, any predicate V and uniformly random $\Theta \in \mathcal{Y}^n$:

$$\begin{aligned} &\Pr\left[\boldsymbol{x} = \boldsymbol{x}' \wedge V(\boldsymbol{x}, \Theta, z) : (\pi, \pi(\boldsymbol{x}), z) \leftarrow \langle \mathcal{S}^{\mathcal{A}}, \pi(\Theta) \rangle\right] \\ \geq\ &\frac{1}{(2q+1)^{2n}} \Pr\left[\boldsymbol{x} = \boldsymbol{x}' \wedge V(\boldsymbol{x}, \mathsf{H}(\boldsymbol{x}), z) : (\boldsymbol{x}, z) \leftarrow \mathcal{A}^{\mathsf{H}}\right]. \end{aligned}$$

Resampling lemma. The resampling lemma [38] says that if an input with sufficiently high min-entropy is chosen and the corresponding output is reprogrammed to a fresh, uniformly random value (it is *resampled*), a polynomial-query distinguisher can detect the reprogramming with a small probability only. For a precise statement, let REPRO_0 and REPRO_1 refer to two games, where after an initial learning phase with q queries, in REPRO_1 the random oracle is reprogrammed as described, while it is left unchanged in REPRO_0.

Lemma 4 ([38, Proposition 2]). *Let $\mathcal{X}_1, \mathcal{X}_2, \mathcal{X}'$ and $\mathcal{Y}$ be finite sets, and let p be a distribution on $\mathcal{X}_1 \times \mathcal{X}'$. Let Dist be any distinguisher, issuing q (quantum) queries to O and R reprogramming instructions such that each instruction consists of a value x_2, together with the fixed distribution p. Then*

$$\left|\Pr\left[1 \leftarrow \mathsf{REPRO}_1^{\mathsf{Dist}}\right] - \Pr\left[1 \leftarrow \mathsf{REPRO}_0^{\mathsf{Dist}}\right]\right| \leq \frac{3R}{2}\sqrt{q \cdot p_{\max}}$$

where $p_{\max} := \max_{x_1} p(x_1)$.

Query Complexity bounds in the QROM. For query bounds in the QROM we rely on the technique from [20] which builds on Zhandry's compressed oracle framework [66]. The compressed oracle is a simulation of the quantum oracle of a random function $H : \mathcal{X} \to \mathcal{Y}$. Its internal state of the oracle can be thought of as a (superposition of) lazy-sampling-style databases of input-output pairs. The probability that $\mathcal{A}$ succeeds in a search task can be related to the probability of the database D, obtained by measuring the internal state of the compressed oracle after the interaction with $\mathcal{A}$, satisfying a certain property related to the search task (see Lemma 5 below). We can think of the database D as a partial function from $\mathcal{X}$ to $\mathcal{Y}$. We write $D(x) = \perp$ if D is undefined on x. We denote the set of all possible such databases by $\mathcal{D}$. For $D \in \mathcal{D}$, $x \in \mathcal{X}$ and $y \in \mathcal{Y} \cup \{\perp\}$, we define $D[x \mapsto y]$ by $D[x \mapsto y](x) = y$ and $D[x \mapsto y](x') = D(x')$ for $x' \neq x$. It will suffice for us to use the *(quantum) transition capacity* formalism from [20], so we will not introduce the compressed oracle in detail.

A subset $P \subseteq \mathfrak{D}$ is called a *database property*, we say that $D \in \mathfrak{D}$ *satisfies* P. The complement is denoted by $\neg P = \mathfrak{D} \setminus P$. For a database property P, a database D and an input $x \in \mathcal{X}$, we define the corresponding local property as

$$P|_{D|^x} = \{y \in \mathcal{Y} | D[x \mapsto y] \in P\}.$$

The maximal probability of a database D satisfying P, when D is obtained by measuring the internal state of the compressed oracle after interaction with $\mathcal{A}$, maximized over all q-query algorithms $\mathcal{A}$, is defined as

$$[\![\bot \Rightarrow^q P]\!] := \max_{\mathcal{A}} \sqrt{\Pr[D \in P]}.$$

Its square is an upper bound on the probability of $\mathcal{A}$ producing such a database.

Lemma 5.6 in [20] shows that the probability of fulfilling a database property after q queries can be bounded using the quantum transition capacity. More precisely, a sequence $P_0, P_1, \ldots, P_q$ with $\neg P_0 = \{\bot\}$ and $P_q = P$ yields

$$[\![\bot \Rightarrow^q P]\!] \leq \sum_{s=0}^{q-1} [\![\neg P_s \to P_{s+1}]\!]$$

where each term is the quantum transition capacity between the databases. We do not define the transition capacity as we will only use it as a formal tool. The transition capacity can be bounded using the following.

Theorem 2 ([28, Theorem 2.4]). *Let P and P' be database properties with trivial intersection, i.e. $P \cap P' = \emptyset$, and for every $D \in \mathcal{D}$ and $x \in \mathcal{X}$ let*

$$L^{x,D} := \begin{cases} P|_{D|^x} & \text{if } \bot \in P'|_{D|^x} \\ P'|_{D|^x} & \text{if } \bot \in P|_{D|^x} \end{cases}$$

with $L^{x,D}$ being either of the two if $\bot \notin P|_{D|^x} \cup P'|_{D|^x}$. Then

$$[\![P \to P']\!] \leq \max_{x,D} \sqrt{10 \Pr[U \in L^{x,D}]}, \tag{1}$$

where U is uniform over $\mathcal{Y}$, and the maximization can be restricted to $D \in \mathcal{D}$ and $x \in \mathcal{X}$ for which both $P|_{D|^x}$ and $P'|_{D|^x}$ are non-empty.

The *fundamental lemma* of the compressed oracle relates the database to the knowledge of a q-query adversary, which yields query bounds.

Lemma 5 ([28, Lemma 2.6]). *Let $\mathcal{A}$ be an oracle quantum algorithm that outputs $\boldsymbol{x} = (x_1, \ldots, x_\ell) \in \mathcal{X}^\ell$ and $z \in \mathcal{Z}$. Let $\tilde{\mathcal{A}}$ be the oracle quantum algorithm that runs $\mathcal{A}$, makes ℓ classical queries on the outputs x_i to obtain $\boldsymbol{y} = \mathsf{H}(\boldsymbol{x})$, and then outputs $(\boldsymbol{x}, \boldsymbol{y}, z)$. When $\tilde{\mathcal{A}}$ interacts with the compressed oracle instead, and at the end D is obtained by measuring the internal state of the compressed oracle, then, conditioned on $\tilde{\mathcal{A}}$'s output $(\boldsymbol{x}, \boldsymbol{y}, z)$,*

$$\Pr[\boldsymbol{y} = D(\boldsymbol{x}) \mid (\boldsymbol{x}, \boldsymbol{y}, z)] \geq 1 - 2\ell|\mathcal{Y}|^{-1}.$$

2.4 Deutsch-Jozsa Algorithm

Let $f : \{0,1\}^n \to \{0,1\}$ be an n-bit function. As a quantum circuit, the Deutsch-Jozsa algorithm [24] can be seen as a measurement in the $+/-$ basis on the first n qubits of the state $U_f |+\rangle^{\otimes n} |-\rangle$, where $U_f |x\rangle |y\rangle := |x\rangle |y \oplus f(x)\rangle$, $x \in \{0,1\}^n$ and $y \in \{0,1\}$. The amplitude of the state $|+\rangle^{\otimes n}$ is $Z_f/2^n$ where

$$Z_f := \sum_{x \in \{0,1\}^n} (-1)^{f(x)}. \tag{2}$$

3 Quantum Oracle Distribution Switching

We define an oracle O_{f_P} for the i.i.d function $f_P : \mathcal{X} \to \mathcal{Y}$ where $f_P(x) \sim P$ for all $x \in \mathcal{X}$ independently and P is a probability distribution on $\mathcal{Y}$.

Consider an algorithm $\mathcal{A}$ interacting with an oracle O_{f_P} or an oracle O_{f_Q}. We characterize $\mathcal{A}$'s output behavior based on the number of quantum queries q, as well as the statistical distance and Rényi divergence of P and Q.

The distinguishing advantage of an algorithm $\mathcal{A}$ that can make at most q classical oracle queries is bounded as

$$\left|\Pr\left[1 \leftarrow \mathcal{A}^{f_P}\right] - \Pr\left[1 \leftarrow \mathcal{A}^{f_Q}\right]\right| \leq q\Delta(P,Q) \tag{3}$$

and for the Rényi divergence assuming $\mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q)$

$$\Pr\left[1 \leftarrow \mathcal{A}^{f_P}\right] \leq \left(R_\alpha(P\|Q)^q \Pr\left[1 \leftarrow \mathcal{A}^{f_Q}\right]\right)^{\frac{\alpha-1}{\alpha}}. \tag{4}$$

3.1 Oracle Switching using Statistical Distance

In this section, let $\epsilon = \Delta(P,Q)$. Boneh et al. [12, Lemma 3] show that the output distributions of $\mathcal{A}$ have statistical distance $\mathcal{O}\left(q^2\sqrt{\epsilon}\right)$ when P or Q is the uniform distribution. This was improved to $\mathcal{O}\left(q^{1.5}\sqrt{\epsilon}\right)$ in [64, Section 7.2]. Using the compressed oracle technique, we show that the density matrices of $\mathcal{A}$'s outputs have trace distance at most $8q\,\mathrm{dh}(P,Q) \leq 8q\sqrt{\epsilon}$.

Theorem 3. *An algorithm $\mathcal{A}$ making q quantum queries to O_{f_P} or O_{f_Q} has the same output distribution, up to statistical distance $8q\,\mathrm{dh}(P,Q) \leq 8q\sqrt{\epsilon}$.*

The proof can be found in Appendix B.1. Therefore, the quantum analog of Eq. (3) only incurs a square-root loss in the statistical distance.

Complexity Theory Literature. [7] characterizes the query complexity of distinguishing P from Q as $\Theta(1/\,\mathrm{dh}(P,Q))$, where dh denotes the Hellinger distance. Theorem 3 is expressed in the same quantity and yields the matching lower bound $q = \Omega(1/\,\mathrm{dh}(P,Q))$ for constant distinguishing advantage. It is, however, obtained by a different route. The bound of [7] is proven via the adversary

method, which characterizes the query complexity at constant success probability and does not provide the explicit constants needed for concrete parameters, whereas we bound the trace distance of the output states at every advantage.

One can use the result to get a per-query statement, but at a square-root loss. Consider a distinguisher $\mathcal{A}$ with advantage η, that outputs 1 with probability $1/2+\eta/2$ when interacting with O_{f_P} and with probability $1/2-\eta/2$ when interacting with O_{f_Q}. By $\Theta(1/\eta^2)$ independent repetitions (domain-separating the repetitions) and a majority vote, the success probability is boosted to $2/3$. A single repetition uses q queries, so $\mathcal{O}\big(q/\eta^2\big)$ queries in total, while [7] requires $\Omega(1/\operatorname{dh}(P,Q))$ queries for constant success probability. Hence, $q/\eta^2 = \Omega(1/\operatorname{dh}(P,Q))$, i.e.,

$$\eta = \mathcal{O}\Big(\sqrt{q \cdot \operatorname{dh}(P,Q)}\Big).$$

This is weaker than Theorem 3 for all $q = o(1/\operatorname{dh}(P,Q))$, and the two coincide only at $q = \Theta(1/\operatorname{dh}(P,Q))$, where both give constant advantage.

3.2 Oracle Switching using Rényi Divergence

The Rényi divergence measures how close two probability distributions are *multiplicatively*. The post-processing and probability-preservation properties (see Lemma 1) facilitate relative-error bounds on $\mathcal{A}$'s outputs when restricted to classical queries. We explore how this behavior generalizes to quantum access.

Replacing the entire Distribution Fails Eq. (4) does not translate to the quantum oracle access setting, i.e., a purely multiplicative bound is impossible. Any such bound has a multiplicative loss depending only on P and Q (in particular, on their Rényi divergences), and is therefore independent of n, the input size for f_P and f_Q. We show that there is always a choice of n for which the quotient of success probabilities exceeds that loss.

Theorem 4. *For all distributions $P \neq Q$ on $\mathcal{Y}$ with $\operatorname{Supp}(P) \subseteq \operatorname{Supp}(Q)$ and fixed $\gamma > 0$, there exists an algorithm $\mathcal{A}$ making 2 quantum queries such that*

$$\lim_{n\to\infty} \Pr\big[1 \leftarrow \mathcal{A}^{f_P}\big] / \Pr\big[1 \leftarrow \mathcal{A}^{f_Q}\big]^{\gamma} = \infty.$$

Proof. The high-level strategy is as follows. We define a classical deterministic algorithm that takes a bit string as an input, makes 2 queries to the oracle given to the distinguisher and outputs one bit, such that the boolean function it implements behaves differently for oracle f_P and oracle f_Q. For oracle f_Q, the boolean function is *balanced in expectation* over the choice of f_Q, meaning that in expectation, the outputs 0 and 1 occur the same number of times. For f_P, on the other hand, it is not. The Deutsch-Jozsa algorithm therefore almost never reports "constant" for f_Q, with probability inversely proportional to the domain size 2^n, whereas for f_P it does so with constant probability. The quotient of the two therefore diverges, and does so even when the f_Q probability is raised to any fixed power $\gamma > 0$.

Fix $\mathcal{X} = \{0,1\}^n$ and define the point of maximal probability quotient $y_0 := \arg\max_{y \in \mathrm{Supp}(P)} \frac{P(y)}{Q(y)}$, for which $Q(y_0) \in (0,1)$ and $P(y_0)/Q(y_0) > 1$ due to $P \neq Q$ and $\mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q)$. From any function $f : \mathcal{X} \to \mathcal{Y}$, we construct a function with a binary output, $g[f] : \mathcal{X} \to \{0,1\}$, where $g[f](x) = 0$ if $f(x) = y_0$ and $g[f](x) = 1$ otherwise. Since this function is not necessarily balanced, we define a padded function $h[f] : \{0,1\}^{1+n} \to \{0,1\}$ increasing the size of the input domain, where $h[f](b\|x) = g[f](x)$ if $b = 0$. For $b = 1$, $h[f]$ is defined to be 0 for exactly $N_0 := \lfloor 2^n(1 - Q(y_0)) \rfloor$ elements (say for concreteness the first N_0 inputs in lexicographic order). The $b = 1$ half is thus chosen so that its bias cancels the expected bias of $g[f_Q]$, up to the rounding error $\delta := 2^n(1 - Q(y_0)) - N_0 \in [0,1)$. Each evaluation of $h[f]$ uses two queries to f: one to compute the comparison $f(x) \stackrel{?}{=} y_0$ into an ancilla, and one to uncompute it after the phase has been applied. Since the uncomputation restores the ancilla to $|0\rangle$, the analysis below is unaffected. After running the DJ algorithm with f (post-processed to $h[f]$), the probability of measuring $|+\rangle^{\otimes(1+n)}$ is

$$\Pr\left[1 \leftarrow \mathcal{A}^f\right] = \mathbb{E}_f\left[\frac{Z_{h[f]}^2}{2^{2(1+n)}}\right] = \frac{\mathbb{E}_f\left[Z_{h[f]}^2\right]}{2^{2(1+n)}} = \frac{\mathbb{E}_f\left[Z_{h[f]}\right]^2 + \mathrm{Var}_f\left(Z_{h[f]}\right)}{2^{2(1+n)}}$$

using the notation from Eq. (2). Since the $b = 1$ half contributes the constant $2N_0 - 2^n = -\mathbb{E}_{f_Q}[Z_{g[f_Q]}] - 2\delta$, we have

$$\mathbb{E}_f[Z_{h[f]}] = \mathbb{E}_f[Z_{g[f]}] - \mathbb{E}_{f_Q}[Z_{g[f_Q]}] - 2\delta,$$

so that $\mathbb{E}_{f_Q}\left[Z_{h[f_Q]}\right] = -2\delta$ and $\mathbb{E}_{f_P}[Z_{h[f_P]}] = 2^{1+n}(P(y_0) - Q(y_0)) - 2\delta$. Since f_P and f_Q assign each input an independent sample, the random variables $(-1)^{g[f](x)}$ for $x \in \{0,1\}^n$ are i.i.d., and the $b = 1$ half of $h[f]$ is deterministic. Hence, the variance is additive:

$$\begin{aligned}\mathrm{Var}_f\left(Z_{h[f]}\right) &= \sum_{x \in \{0,1\}^n} \mathrm{Var}_f\left((-1)^{g[f](x)}\right) \\ &= \sum_{x \in \{0,1\}^n} 4\Pr[g[f](x) = 0]\left(1 - \Pr[g[f](x) = 0]\right) = 2^{2+n} v_f\end{aligned}$$

where $v_f = P(y_0)(1 - P(y_0))$ if $f = f_P$ and $Q(y_0)(1 - Q(y_0))$ if $f = f_Q$. Using $\delta < 1$, the rounding contributes $\mathcal{O}\left(2^{-2n}\right)$ to $\Pr\left[1 \leftarrow \mathcal{A}^{f_Q}\right]$, and $\mathcal{O}(2^{-n})$ to $\Pr\left[1 \leftarrow \mathcal{A}^{f_P}\right]$ via the cross term $-2\delta\left(P(y_0) - Q(y_0)\right)2^{-n}$. Since $Q(y_0) \in (0,1)$, the denominator is nonzero, and thus

$$\frac{\Pr\left[1 \leftarrow \mathcal{A}^{f_P}\right]}{\Pr[1 \leftarrow \mathcal{A}^{f_Q}]^\gamma} = \frac{\left(P(y_0) - Q(y_0)\right)^2 + \mathcal{O}(2^{-n})}{\left(2^{-n}Q(y_0)(1 - Q(y_0)) + \mathcal{O}(2^{-2n})\right)^\gamma} = \Theta(2^{n\gamma})$$

tends to ∞ in n for any fixed choice of γ. □

Replacing some Positions using Small-Range Distributions The counterexample makes crucial use of the fact that the multiplicative bounds should hold for any absolute value for the probabilities. Adding an additive error term could be sufficient to achieve a meaningful result for the Rényi divergence, even if the classical properties do not translate. We explore this approach by considering small-range distributions from [64]. In short, our approach relies on first sampling r independent values according to a distribution P on $\mathcal{Y}$. These are used to simulate the oracle O_{f_P}. It is indistinguishable for $\mathcal{A}$ up to an error term of $\mathcal{O}(q^3/r)$ using [64]. Next, the distribution is changed from P to Q. Now we can use the properties of the Rényi divergence for r replacements.

Proposition 1. *Let $\alpha \in (1,\infty], r \in \mathbb{N}_{\geq 1}$ and P, Q be classical distributions over $\mathcal{Y}$ with $\mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q)$. For any $\mathcal{A}$ making at most q quantum queries*

$$\Pr\left[1 \leftarrow \mathcal{A}^{f_P}\right] \leq \left(R_\alpha(P\|Q)^r \left(\Pr\left[1 \leftarrow \mathcal{A}^{f_Q}\right] + \frac{\ell(q)}{r}\right)\right)^{\frac{\alpha-1}{\alpha}} + \frac{\ell(q)}{r},$$

where $\ell(q) = \pi^2(2q)^3/3 < 27q^3$.

Remark 1. The result is partially for exposition and study of how Rényi divergence can be used by accepting a small additive loss while still switching between the entire oracles. This technique could be useful in some situations. Statistical distance arguments, however, often yield tighter bounds. This is because selecting r such that the error is negligible implies a small statistical distance, sufficient to apply the results from Theorem 3. The full discussion on applicability can be found in Appendix B.2.

3.3 Reprogramming with Different Distributions

Though possible for the statistical distance, it is impossible to use the Rényi divergence for PQ security as in the classical setting. For a fixed oracle, replacing all underlying values with another distribution is infeasible when using the Rényi divergence. However, it is possible as shown in Section 3.2, but at the cost of an additive error.

In many applications, it turns out to be sufficient to *reprogram* a quantum oracle for an i.i.d. function with output distribution P using outputs sampled according to Q, *for a number of randomly sampled inputs*. The adaptive reprogramming approach can be modified to facilitate this (see Lemma 4). Let $\mathsf{REPRO}_{1,Q}$ refer to a game, modified from REPRO_1, where the output is reprogrammed to a fresh, random value sampled from a distribution Q.

Lemma 6 (Adaptive Reprogramming with Distribution Switching). *Let $\mathcal{X}_1, \mathcal{X}_2, \mathcal{X}'$ and $\mathcal{Y}$ be finite sets, and let p be a distribution on $\mathcal{X}_1 \times \mathcal{X}'$. Let Q be a distribution over $\mathcal{Y}$ and Dist be any distinguisher, issuing q (quantum) queries to O and R reprogramming instructions such that each instruction consists of a*

value x_2, together with the fixed distribution p. Then

$$\left|\Pr\left[1 \leftarrow \mathsf{REPRO}_0^{\mathsf{Dist}}\right] - \Pr\left[1 \leftarrow \mathsf{REPRO}_{1,Q}^{\mathsf{Dist}}\right]\right| \leq R \cdot \Delta(Q, \mathcal{U}(\mathcal{Y})) + \delta_{repr}$$
$$\Pr\left[1 \leftarrow \mathsf{REPRO}_0^{\mathsf{Dist}}\right] \leq \left(R_\alpha(\mathcal{U}(\mathcal{Y})\|Q)^R \Pr\left[1 \leftarrow \mathsf{REPRO}_{1,Q}^{\mathsf{Dist}}\right]\right)^{\frac{\alpha-1}{\alpha}} + \delta_{repr}$$

where $\delta_{repr} = \frac{3R}{2}\sqrt{q \cdot p_{\max}}$ *with* $p_{\max} := \max_{x_1} p(x_1)$ *and* $\alpha \in (1, \infty]$.

The lemma is an immediate consequence of Lemma 4 when applying the statistical and Rényi divergence properties to the classical samples used in reprogramming. A proof, outlining the equations, can be found in Appendix B.3. The result can be extended to two arbitrary underlying distributions.

4 A Closer Look at Falcon in the QROM

Our analysis in Section 3 can be applied to the NIST candidate Falcon [33]. Falcon's security in the QROM has not been rigorously established in existing literature. [33] argues that Falcon is secure because it uses the GPV-framework, which was proven using a history-free reduction in [12]. However, a closer look reveals that this line of argument is problematic. The QROM proof of [12] requires quantum oracle indistinguishability, which is proven via statistical distance arguments (Theorem 3 could be used to get tighter arguments than those in [12]). The parameters of Falcon, however, are set using ROM security considerations via Rényi divergence arguments[2]. Although it is possible to use the Rényi divergence to get bounds on the statistical distance using Pinsker's inequality, the bounds that are derived this way are impractical ([34] shows that Falcon has statistical distance $\approx 2^{-34}$). This renders [12] infeasible for Falcon.

Classically, the Rényi divergence yields a multiplicative bound (see the ROM proof for Falcon in [34]). This directly allows using the data processing inequality (see Lemma 1). A straightforward generalization of this strategy to the QROM is impossible as we have shown in Section 3.2. Tolerating additional error terms, we could use Proposition 1, but this implies a small statistical distance, rendering this approach unusable for Falcon's parameters as well.

As we have sufficient entropy in the hash inputs, we can use our results from Section 3.3 to only reprogram those queries where a call on $\mathsf{PreSmp}_{\mathsf{NTRU}}$ is performed, enabling use of the Rényi divergence. To obtain a tighter reduction, Falcon uses a multi-target assumption (t-$\mathcal{R}$-NTRU-ISIS) when showing EUF-CMA security. In the QROM, the only technique we are aware of that supports reduction to this assumption is small-range distributions [64]. However, the number of targets must increase significantly: classically $t = q_{\mathsf{H}} + 1$ suffices with no additive error in ϵ, whereas in the QROM an additive error of η in ϵ costs $t > 27(q_{\mathsf{H}} + C_s + 1)^3/\eta$ (see the discussion below Theorem 5).

[2] It also requires considering aborts, we address this separately in the proof.

Although [34, Section 5.1] does make the assumption that t-$\mathcal{R}$-NTRU-ISIS and t-$\mathcal{R}$-NTRU-SPISIS are as hard as $\mathcal{R}$-NTRU-SIS independent of t, we believe this assumption warrants additional scrutiny due to the significantly larger t. We also need to work with a slightly different version t-$\mathcal{R}$-NTRU-SPISIS* of the assumption from [34, Definition 10] (t-$\mathcal{R}$-NTRU-SPISIS), as conditioning the challenges on the norm bound, as they do, is incompatible with applying the Rényi arguments per sample. This is only needed when proving SUF-CMA, and not needed for EUF-CMA (see Remark 3). [34] also acknowledges that more efficient attacks against t-$\mathcal{R}$-NTRU-ISIS for large t may exist [10]. For this reason, we consider the alternative reduction to single-target $\mathcal{R}$-NTRU-ISIS via measure-and-reprogram (Theorem 1). This is simpler and only incurs the standard square root loss compared to the analogous classical reduction. On the other hand, this strategy yields less tight reductions compared to the reduction to multi-target $\mathcal{R}$-NTRU-ISIS.

The details of the proofs and Falcon can be found in Appendix A.6 and Appendix E.1. The version of Falcon we analyze is CoreFalcon$^+$ [34].

Theorem 5 (SUF-CMA of CoreFalcon$^+$). *For any adversary $\mathcal{A}$ against the* SUF-CMA *security of CoreFalcon$^+$, making at most q_s signing queries and q_{H} QROM queries, there exists an adversary $\mathcal{B}$ against* C$_s$-$\mathcal{R}$-NTRU-SPISIS**, an adversary $\mathcal{C}$ against* $\mathcal{R}$-NTRU-ISIS *and an adversary $\mathcal{D}$ against* t-$\mathcal{R}$-NTRU-ISIS*, s.t. for all $C_s, t \in \mathbb{N}^{\geq 1}, \alpha_u, \alpha_p \in \mathbb{R}^{>1}$:*

$$\mathsf{Adv}^{\text{SUF-CMA}}_{\mathcal{A},\textsc{CoreFalcon}^+,q_s}(\lambda) \leq \left(r_u^{C_s} \left(r_p^{C_s} \left(\epsilon + \epsilon' \right) \right)^{\frac{\alpha_p - 1}{\alpha_p}} \right)^{\frac{\alpha_u - 1}{\alpha_u}} + \delta$$

where ϵ and the runtimes of $\mathcal{B}$ and $\mathcal{D}$ are bounded in Table 1 and the other parameters, including δ as a function of C_s which is in a trade-off relationship with t via epsilon, are defined in Table 2.

The two reductions to single- and multi-target assumptions have different advantages and drawbacks. The single-target reduction via measure-and-reprogram has the advantage of only suffering the standard Grover-type loss compared to the corresponding classical reduction. It is, however, not tight. The multi-target reduction has the same success probability tightness advantage as its classical analgue. In contrast to the classical reduction, however, the use of the small-range distribution technique (intuitively necessitated by the fact that a single superposition query can explore all QRO outputs) yields a reduction runtime polynomial in the runtime and the inverse success probability of the adversary (seen by balancing the two summands in the bound on ϵ in Table 1). Additionally, the additive error δ that results from the use of reprogramming and thus has no analogue in the classical reduction required a relatively large salt length t.

Table 1. Trade-offs between runtime and assumptions for ϵ in Theorem 5 when reducing to CoreFalcon$^+$. C_s is a parameter of the reduction that can be set to $\Theta(q_s/p)$, where $1-p$ is the rejection probability of the rejection sampling routine. In [34, Corollary 3 to 6] it was for example set as $C_s < 2q_s$.

Assumption	Reduction Runtime	Bound on ϵ
single-target	$\mathsf{poly}(\lambda, C_s, q_\mathsf{H})$	$(2q_\mathsf{H}+2C_s+3)^2\mathsf{Adv}^{\mathcal{R}\text{-NTRU-ISIS}}_{\mathcal{C},1,q,\alpha,\beta}(\lambda)$
multi-target	$\mathsf{poly}(\lambda, C_s, t)$	$\mathsf{Adv}^{\mathrm{t}\text{-}\mathcal{R}\text{-NTRU-ISIS}}_{\mathcal{D},1,q,\alpha,\beta}(\lambda) + \frac{27(q_\mathsf{H}+C_s+1)^3}{t}$

Table 2. Parameters from Theorem 5 where $\mathcal{Q}_{\boldsymbol{h}}$ is the distribution of $\boldsymbol{h}\boldsymbol{s}_1+\boldsymbol{s}_2 \bmod q$ with $\boldsymbol{s}_1, \boldsymbol{s}_2 \leftarrow \mathcal{D}_{\mathcal{R},s}$, and where $\boldsymbol{\Lambda}_{\boldsymbol{c}} := \boldsymbol{\Lambda}(\boldsymbol{B}) + (\boldsymbol{0}, \boldsymbol{c})$ denotes the coset of the NTRU lattice containing the preimages of $\boldsymbol{c}$, the output of $\mathsf{PreSmp}_{\mathsf{NTRU}}$ being read in the order $(\boldsymbol{s}_2, \boldsymbol{s}_1)$.

	Definition
ϵ'	$\mathsf{Adv}^{\mathrm{C_s}\text{-}\mathcal{R}\text{-NTRU-SPISIS}^*}_{\mathcal{B},q,\alpha,\beta}(\lambda) + \dfrac{10(C_s+q_\mathsf{H}+1)^2 C_s}{\lvert\mathcal{R}_q\rvert}$
δ	$\dfrac{3C_s}{2}\sqrt{\dfrac{C_s+q_\mathsf{H}+1}{2^k}} + \displaystyle\sum_{i=0}^{q_s}\binom{C_s}{i}(1-p)^{C_s-i}(p)^i$
r_u	$\max_{(\boldsymbol{h},\cdot,\cdot)\in\mathrm{Supp}(\mathsf{TpdGen}_{\mathsf{NTRU}})} R_{\alpha_u}(\mathcal{U}(\mathcal{R}_q)\lVert\mathcal{Q}_{\boldsymbol{h}})$
r_p	$\max_{(\boldsymbol{h},\boldsymbol{f},\boldsymbol{g})\in\mathrm{Supp}(\mathsf{TpdGen}_{\mathsf{NTRU}}),\, \boldsymbol{c}\in\mathcal{R}_q} R_{\alpha_p}(\mathsf{PreSmp}_{\mathsf{NTRU}}(\boldsymbol{B}_{\boldsymbol{f},\boldsymbol{g}}, s, \boldsymbol{c})\lVert\mathcal{D}_{\boldsymbol{\Lambda}_{\boldsymbol{c}},s})$
p	$\min_{(\boldsymbol{h},\boldsymbol{f},\boldsymbol{g})\in\mathrm{Supp}(\mathsf{TpdGen}_{\mathsf{NTRU}}),\, \boldsymbol{c}\in\mathcal{R}_q} \Pr\left[\substack{\lVert(\boldsymbol{s}_1,\boldsymbol{s}_2)\rVert_2\le\beta: \\ (\boldsymbol{s}_1,\boldsymbol{s}_2)\leftarrow\mathsf{PreSmp}_{\mathsf{NTRU}}(\boldsymbol{B}_{\boldsymbol{f},\boldsymbol{g}},s,\boldsymbol{c})}\right]$

5 Ring Signature from Ring Preimage Sampleable Functions

We analyse the QROM security of ring-trapdoor-based ring signatures. We first formulate a notion of RTDFs to capture these ring signatures.

Definition 1 (Ring Trapdoor Function). *A ring trapdoor function (RTDF) is a quadruple of ppt algorithms* $(\mathsf{Stp}, \mathsf{TpdGen}, f, \mathsf{SamplePre})$ *defined as follows:*

- $\mathsf{Stp}(1^\lambda) \to \mathsf{p}$*: The setup algorithm takes in the security parameter* 1^λ *and outputs public parameters* p*, that also define the maximal ring size* κ*.*
- $\mathsf{TpdGen}(\mathsf{p}) \to (a, t)$*: The trapdoor generation algorithm takes the public parameter* p *as input and produces a pair* (a, t) *with public value* a *and trapdoor* t*. Every ring* $\rho = \{a_1, \ldots, a_N\}$ *with* $N \leq \kappa$ *defines a domain* D_ρ*, a finite range* R_ρ *and an efficient function* $f_\rho : \mathsf{D}_\rho \to \mathsf{R}_\rho$*.*
- $\mathsf{SamplePre}(\rho, t_i, r) \to d$*: The presampling algorithm takes in a ring* ρ*, a trapdoor* t *and an element* $r \in \mathsf{R}_\rho$*. The ring must satisfy the size bound, i.e.,* $N = |\rho| \leq \kappa$*, and* $\exists i \in [N]$ *such that* $\mu(t) = \rho_i$[3]*. It outputs a* $d \in \mathsf{D}_\rho$*.*

A RTDF is $\delta(\kappa)$*-correct if for any* $\mathsf{p} \leftarrow \mathsf{Stp}(1^\lambda)$*, all* $N \leq \kappa$*, any* $\{(a_j, t_j)\}_{j\in[N]} \subseteq \mathrm{Supp}(\mathsf{TpdGen}(\mathsf{p}))$ *defining* $\rho = \{a_1, \ldots, a_N\}$ *and any* $r \in \mathsf{R}_\rho$

$$\Pr[f_\rho(\mathsf{SamplePre}(\rho, t_j, r)) \neq r] \leq \delta(\kappa).$$

We introduce properties for RTDFs similar to those of PSFs [36], in particular, the existence of an efficient domain sampling algorithm $\mathsf{SampleDom}(\rho)$ that takes a ring ρ with $|\rho| \leq \kappa$ as input and outputs $d \in \mathsf{D}_\rho$ satisfying the properties below. If a RTDF satisfies all of these properties, we also refer to it as a RPSF. Let $\mathsf{p} \leftarrow \mathsf{Stp}(1^\lambda), \{(a_i, t_i)\}_{i\in[\kappa]} \leftarrow \mathsf{TpdGen}(\mathsf{p})$ and $\rho = \{a_i\}_{i\in[\kappa]}$.

1. *Domain sampling with uniform output*: Define $\epsilon^{dom}_{\mathsf{RPSF}}$ and $\delta^{\alpha\text{-}dom}_{\mathsf{RPSF}}$ such that for any ρ' with $\rho \cap \rho' \neq \emptyset$ it holds that

$$\Delta(D, \mathcal{U}(\mathsf{R}_{\rho'})) \leq \epsilon^{dom}_{\mathsf{RPSF}} \quad \text{and} \quad R_\alpha(\mathcal{U}(\mathsf{R}_{\rho'}) \| D) \leq \delta^{\alpha\text{-}dom}_{\mathsf{RPSF}}$$

   where $D := f_{\rho'}(\mathsf{SampleDom}(\rho'))$.
2. *Preimage sampling is not detectable*: Define three upper bounds $\epsilon^{pre}_{\mathsf{RPSF}}, \delta^{\alpha\text{-}pre}_{\mathsf{RPSF}}$ and $\epsilon^{KL\text{-}pre}_{\mathsf{RPSF}}$ such that for every ρ' with $\rho \cap \rho' \neq \emptyset$ and $\forall i \in [\kappa]$ it holds that

$$\Delta(D, Q) \leq \epsilon^{pre}_{\mathsf{RPSF}}, \quad R_\alpha(D \| Q) \leq \delta^{\alpha\text{-}pre}_{\mathsf{RPSF}} \quad \text{and} \quad D_{KL}(D \| Q) \leq \epsilon^{KL\text{-}pre}_{\mathsf{RPSF}}$$

   where $r \leftarrow \mathsf{R}_{\rho'}$, $D := \mathsf{SamplePre}(\rho', t_i, r)$ and Q is the distribution of $d \leftarrow \mathsf{SampleDom}(\rho')$ conditioned on $f_{\rho'}(d) = r$ for $\alpha \in (1, \infty]$.

[3] We assume (w.l.o.g.) that there exists a one-way function μ such that $\mu(t) = a$ for all $(a, t) \in \mathrm{Supp}(\mathsf{TpdGen})$.

3. *One-wayness*: For a ppt algorithm $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$, define the one-wayness

$$\mathsf{Adv}^{\mathrm{OW}}_{\mathcal{A},\mathsf{RPSF}}(\lambda) = \Pr\left[f_{\rho'}(d) = r \wedge \rho' \subseteq \rho : \begin{array}{l} (\mathsf{state}, \rho') \leftarrow \mathcal{A}_1(1^\lambda, \rho) \\ \wedge r \leftarrow \mathsf{R}_{\rho'} \\ \wedge d \leftarrow \mathcal{A}_2(\mathsf{state}, r) \end{array} \right].$$

4. *Preimage min-entropy*: The min-entropy of $\mathsf{SampleDom}(\rho')$ conditioned on $f_{\rho'}(d) = r$ for $\rho' \subseteq \rho$ is at least $\beta(\lambda)$ if for every $r \in \mathsf{R}_{\rho'}$ and every $D \in \mathsf{D}_{\rho'}$

$$\Pr[D = d : d \leftarrow \mathsf{SampleDom}(\rho') \mid f_{\rho'}(d) = r] \leq 2^{-\beta(\lambda)}.$$

5. *Collision-resistance*: For any ppt algorithm $\mathcal{A}$, define the collision resistance

$$\mathsf{Adv}^{\mathrm{COL}}_{\mathcal{A},\mathsf{RPSF}}(\lambda) = \Pr\left[f_{\rho'}(d_1) = f_{\rho'}(d_2) \wedge d_1 \neq d_2 \wedge \rho' \subseteq \rho : (d_1, d_2, \rho') \leftarrow \mathcal{A}(1^\lambda, \rho) \right].$$

We construct a generic ring signature from a RPSF as in Fig. 1. Each party generates a RPSF trapdoor (sk) and public value (pk). To sign a message for a fixed ring, hash the ring and the message together with a salt salt to compute a target h. We then use the RPSFs presample algorithm to get a preimage. The signature is the preimage and the salt. A signature is verified by recomputing the target and checking that the signature maps to the target under f.

$\mathsf{Sign}(\mathsf{sk}, \rho, m)$	$\mathsf{Vf}(\rho, m, (\mathsf{salt}, \sigma))$
1 : $\mathsf{salt} \leftarrow \{0,1\}^k$	1 : $h \leftarrow \mathsf{H}(\rho, \mathsf{salt}, m)$
2 : $h \leftarrow \mathsf{H}(\rho, \mathsf{salt}, m)$	2 : **return** $\mathsf{RPSF}.f_\rho(\sigma) = h$
3 : $\sigma \leftarrow \mathsf{RPSF.SamplePre}(\rho, \mathsf{sk}, h)$	
4 : **return** (salt, σ)	

Fig. 1. Generic construction of a ring signature from a RPSF with k bits of salt. The key generation and setup algorithm of the ring signature are identical to those of RPSF.

5.1 Unforgeability via Adaptive Reprogramming

The target for the preimage sampler in the signing algorithm in Fig. 1 is computed using a hash function. If the hash input has sufficiently high min-entropy in its inputs to H, e.g. $\lambda = k$, we can model the hash function as a quantum-accessible random oracle and use reprogramming (Lemma 6) to simulate the signing oracle for an SUF-CRA attacker (Lemma 6).

Theorem 6. *Let* RPSF *be a $\delta(\kappa)$-correct RPSF with preimage-min entropy $\beta(\lambda)$, and* RSig *be the generic ring signature from Fig. 1. Let $\alpha_1, \alpha_2 \in (1, \infty]$, $N \leq \kappa, k \in \mathbb{N}$ and* H *be modeled in the QROM. For any* SUF-CRA *adversary $\mathcal{A}$ making at most q_s signing and q_H hash queries, there exists an adversary $\mathcal{A}_{\mathrm{COL}}$ against the collision property of the RPSF, and an adversary $\mathcal{A}_{\mathrm{OW}}$ against the one-wayness of the RPSFs s.t. $\mathsf{Adv}^{\text{SUF-CRA}}_{\mathcal{A},\mathsf{RSig},N,q_s}(\lambda)$ is bounded as in Table 3.*

Table 3. Bounds on $\mathsf{Adv}^{\text{SUF-CRA}}_{\mathcal{A},\mathsf{RSig},N,q_s}(\lambda)$ in Theorem 6 when using statistical distance (SD) or Rényi divergence (RD) where C is the minimal number of elements in R_ρ for any honestly generated ring ρ, $\epsilon' = \frac{3q_s}{2}\sqrt{(q_s+q_\mathsf{H}+1)\cdot\frac{1}{2^k}}$ and $\epsilon = \mathsf{Adv}^{\text{COL}}_{\mathcal{A}_{\text{COL}},\mathsf{RPSF}}(\lambda) + 20(q_s+q_\mathsf{H}+1)^2C^{-1}q_s + (2q_s+2q_\mathsf{H}+3)^2\mathsf{Adv}^{\text{OW}}_{\mathcal{A}_{\text{OW}},\mathsf{RPSF}}(\lambda)$.

Domain	Preimage	Bound on $\mathsf{Adv}^{\text{SUF-CRA}}_{\mathcal{A},\mathsf{RSig},N,q_s}(\lambda)$
RD	RD	$\left((\delta^{\alpha_1\text{-}dom}_{\mathsf{RPSF}})^{q_s}\left((\delta^{\alpha_2\text{-}pre}_{\mathsf{RPSF}})^{q_s}\cdot\epsilon\right)^{\frac{\alpha_2-1}{\alpha_2}}\right)^{\frac{\alpha_1-1}{\alpha_1}} + \epsilon'$
SD	RD	$\left((\delta^{\alpha_2\text{-}pre}_{\mathsf{RPSF}})^{q_s}\cdot\epsilon\right)^{\frac{\alpha_2-1}{\alpha_2}} + q_s\epsilon^{dom}_{\mathsf{RPSF}} + \epsilon'$
RD	SD	$\left((\delta^{\alpha_1\text{-}dom}_{\mathsf{RPSF}})^{q_s}\left(\epsilon + q_s\epsilon^{pre}_{\mathsf{RPSF}}\right)\right)^{\frac{\alpha_1-1}{\alpha_1}} + \epsilon'$
SD	SD	$q_s(\epsilon^{dom}_{\mathsf{RPSF}} + \epsilon^{pre}_{\mathsf{RPSF}}) + \epsilon + \epsilon'$

Proof. The correctness is $\delta(\kappa)$ and follows by construction. The strategy for the SUF-CRA proof is as follows. The random oracle is reprogrammed at every signature query. The hash inputs are domain-separated for different rings. We can reprogram the hash value to a fresh uniform random value, which we then in turn can replace by sampling a value σ in the domain using $\mathsf{SampleDom}$ and using $f_\rho(\sigma)$. Now, the signature σ is freshly uniformly sampled and thus does not depend on the private signing keys anymore, so we can transition to the UF-NRA game. Finally, a forgery breaks the one-wayness of the underlying RPSF.

Game_0: This is the original game, so here we have

$$\mathsf{Adv}^{\text{SUF-CRA}}_{\mathcal{A},\mathsf{RSig},N,q_s}(\lambda) = \Pr\left[\mathsf{Game}^{\mathcal{A}}_0\right].$$

$\mathcal{A}$ eventually outputs a forgery $(\rho^*, (\mathsf{salt}^*, \sigma^*), m^*)$, and we assume that the adversary made a classical query on $(\rho^*, \mathsf{salt}^*, m^*)$. This adds one query to q_H.

Game_1: The hash value in signature query i for ρ_i is reprogrammed to $f_{\rho_i}(\sigma_i)$, where $\sigma_i \leftarrow \mathsf{SampleDom}_{\rho_i}(1^\lambda)$. We use Property 1 and Lemma 6 to get

$$\left|\Pr\left[\mathsf{Game}^{\mathcal{A}}_0\right] - \Pr\left[\mathsf{Game}^{\mathcal{A}}_1\right]\right| \leq q_s\cdot\epsilon^{dom}_{\mathsf{RPSF}} + \epsilon'$$

$$\Pr\left[\mathsf{Game}^{\mathcal{A}}_0\right] \leq \left((\delta^{\alpha_1\text{-}dom}_{\mathsf{RPSF}})^{q_s}\Pr\left[\mathsf{Game}^{\mathcal{A}}_1\right]\right)^{\frac{\alpha_1-1}{\alpha_1}} + \epsilon'$$

where $\epsilon' = \frac{3q_s}{2}\sqrt{(q_s+q_\mathsf{H}+1)\cdot p_{\max}}$ and $\rho'\cap\rho\neq\emptyset$. The hash input has at least k bits of min-entropy from salt, so $p_{\max}\leq\frac{1}{2}^k$.

Game_2: Each signature is replaced by σ_i.[4] This is undetected due to Property 2

$$\left|\Pr\left[\mathsf{Game}_1^{\mathcal{A}}\right] - \Pr\left[\mathsf{Game}_2^{\mathcal{A}}\right]\right| \leq q_s \cdot \epsilon^{pre}_{\mathsf{RPSF}}$$

$$\Pr\left[\mathsf{Game}_1^{\mathcal{A}}\right] \leq \left((\delta^{\alpha_2\text{-}pre}_{\mathsf{RPSF}})^{q_s} \Pr\left[\mathsf{Game}_2^{\mathcal{A}}\right]\right)^{\frac{\alpha_2-1}{\alpha_2}}.$$

We reduce $\mathsf{Game}_2^{\mathcal{A}}$ to the properties of the underlying RPSF. If the final forgery $(\rho^*, (\mathsf{salt}^*, \sigma^*), m^*)$ uses the same target as a previous signature query on (m, salt, ρ), they define two preimages for the same target. If the signatures are the same, then $(m^*, \mathsf{salt}^*) \neq (m, \mathsf{salt})$, and $\mathcal{A}$'s forgery query hits one of the targets fixed by the signing queries. Let $T \subseteq \mathsf{R}_\rho$ denote the set of these targets, so $|T| \leq q_s$. The forgery query lies outside the reprogrammed positions, so its value is uniform, and $\Pr[\gamma(U) \in T] \leq q_s \cdot \gamma_{cl}/|\mathcal{Y}|$ for any fixed T. Substituting this for $\Pr\left[U \in L^{x,D}\right] \leq s \cdot \gamma_{cl}/|\mathcal{Y}|$ in the proof of Lemma 8 gives a transition capacity of $\sqrt{10}\sqrt{q_s \cdot \gamma_{cl}/|\mathcal{Y}|}$, which is independent of s and holds at every step. Summing over the $q_s + q_{\mathsf{H}} + 1$ queries and squaring as in the proof of Lemma 9 bounds the probability that the database contains such a query by $10(q_s + q_{\mathsf{H}} + 1)^2 \cdot q_s \cdot \gamma_{cl}/|\mathcal{Y}| \leq 20(q_s + q_{\mathsf{H}} + 1)^2 q_s / C$[5].

If $\mathcal{A}$ uses a different uniform target, we reduce to one-wayness by replacing the hash query from the final forgery by a uniform value using measure-and-reprogram. This incurs a multiplicative loss of $(2q_s + 2q_{\mathsf{H}} + 3)^2$. In total:

$$\Pr\left[\mathsf{Game}_2^{\mathcal{A}}\right] \leq \mathsf{Adv}^{\mathrm{COL}}_{\mathcal{A},\mathsf{RPSF}}(\lambda) + \frac{20(q_s + q_{\mathsf{H}} + 1)^2 q_s}{C} + (2q_s + 2q_{\mathsf{H}} + 3)^2 \mathsf{Adv}^{\mathrm{OW}}_{\mathcal{A},\mathsf{RPSF}}(\lambda).$$

The four combinations of transitions between Game_0 to Game_1 and Game_1 to Game_2 yield the four rows of Table 3. □

Note that removing the salt does not yield a construction provably with UF-CRA or UF-CRA1 security. The salt is in fact required for the proof technique itself, as the adaptive reprogramming needs the additional min-entropy.

Comparison to the ROM. The proof in the ROM is structurally the same. The terms for adaptive reprogramming and QROM collisions are replaced by terms for salt and ROM collisions. Both are only additive, but the salt length k and size of the range can be made smaller. Lastly, the one-wayness challenge can be included by guessing of one of the $q_{\mathsf{H}} + 1$ hash queries and include the challenge in there. This is roughly a multiplicative difference of q_{H} assuming that $q_s \leq q_{\mathsf{H}}$.

[4] This step implicitly contains the correctness properties of RPSF. The simulated signatures are all valid, but the ones generated with SamplePre might not be actual preimages. This difference in distribution is, however, captured within Property 2, as the property already considers that SamplePre might not output correct preimages.

[5] The hash function maps directly into R_ρ, and hash inputs are domain-separated between rings. Note that γ can be chosen such that hashing into $\mathcal{Y}$ and mapping to R_ρ satisfies $\gamma_{cl}/|\mathcal{Y}| \leq 2/C$ whenever $|\mathcal{Y}| \geq C$.

5.2 Unforgeability via History-Free Proofs

If adding a salt is too costly, we can use the history-free reduction technique from [12]. Instead of targeted reprogramming, this technique replaces all hash outputs by the RPSF applied to random domain samples, i.e., it switches to a different i.i.d. function and the results from Theorem 3 and Section 3.2 apply. Hence, we cannot use the Rényi divergence properties in Property 1.

We can only show UF-CRA1 security of the ring signature without the salt.

Theorem 7. *Let* RPSF *be a* $\delta(\kappa)$*-correct RPSF with preimage-min entropy* $\beta(\lambda)$*, and* RSig *be the generic ring signature from Fig. 1. Let* $\alpha \in (1, \infty], N \leq \kappa, k = 0$ *and* H *be modeled in the QROM. For any* UF-CRA1 *adversary* $\mathcal{A}$ *making at most* q_s *signing and* q_{H} *hash queries, there exists an adversary* $\mathcal{A}_{\mathrm{COL}}$ *against the collision property of the RPSF such that* $\mathsf{Adv}^{\text{UF-CRA1}}_{\mathcal{A},\mathsf{RSig},N,q_s}(\lambda)$ *is bounded as in Table 4.*

Table 4. Bounds on $\mathsf{Adv}^{\text{UF-CRA1}}_{\mathcal{A},\mathsf{RSig},N,q_s}(\lambda)$ in Theorem 7 when using statistical distance (SD) or Rényi divergence (RD) where $\epsilon = 8(q_s + q_{\mathsf{H}})(\epsilon^{dom}_{\mathsf{RPSF}})^{1/2}$.

Preimage	Bound on $\mathsf{Adv}^{\text{UF-CRA1}}_{\mathcal{A},\mathsf{RSig},N,q_s}(\lambda)$
SD	$\epsilon + q_s\epsilon^{pre}_{\mathsf{RPSF}} + 2^{-\beta(\lambda)} + \mathsf{Adv}^{\mathrm{COL}}_{\mathcal{A}_{\mathrm{COL}},\mathsf{RPSF}}(\lambda)$
RD	$\epsilon + \left[(\delta^{\alpha\text{-}pre}_{\mathsf{RPSF}})^{q_s}\left(2^{-\beta(\lambda)} + \mathsf{Adv}^{\mathrm{COL}}_{\mathcal{A}_{\mathrm{COL}},\mathsf{RPSF}}(\lambda)\right)\right]^{\frac{\alpha-1}{\alpha}}$

We prove the QROM security of this generic ring signature via a history-free reduction. This includes the history-free simulation of the random oracle and signing oracle, ensuring consistency. History-free reductions were initially introduced in [12] for signature schemes. We extend the concept to ring signatures (see Appendix C). We augment the formalism to support explicit security bounds and use Theorem 3 to bound the detectability of switching the oracle distribution. The adaptations and the formal proof can be found in Appendix C.

The reduction strategy can be sketched as follows. The random oracle is simulated by using a different, private random oracle O_c to generate randomness for $\mathsf{SampleDom}$, using it to "sample" σ in the domain, and then mapping σ to the range using f. Then, Theorem 3 can be applied. The signing queries can be simulated by using O_c and $\mathsf{SampleDom}$, which is undetectable due to the preimage sampling being undetectable. A valid forgery yields a collision for the RPSF with high probability due to the min-entropy of the domain sampling.

5.3 Anonymity

The anonymity proof uses the same ideas as [35, Theorem 1]. The strategy is independent of the modelling of the hash oracle, so it applies to the QROM.

Corollary 1. *Let* RPSF *be a* $\delta(\kappa)$*-correct RPSF, and* RSig *be the generic ring signature from Fig. 1. Let* H *be modeled in the QROM. For any* ANON *adversary* $\mathcal{A}$ *making at most* q_c *challenge queries*

$$\mathsf{Adv}^{\mathrm{ANON}}_{\mathcal{A},\mathsf{RSig},N,q_c}(\lambda) \leq q_c \epsilon^{pre}_{\mathsf{RPSF}}, \tag{5}$$

$$\mathsf{Adv}^{\mathrm{ANON}}_{\mathcal{A},\mathsf{RSig},N,q_c}(\lambda) \leq \sqrt{\left(q_c \epsilon^{KL\text{-}pre}_{\mathsf{RPSF}}\right)/2}. \tag{6}$$

In the anonymity game, we simulate the CHAL oracle using conditional preimage sampling, rendering the signature manifestly independent of the signer and thus the CHAL oracle independent of b. The change in output behavior of $\mathcal{A}$ can be measured using the statistical distance or the Kullback-Leibler divergence.

Proof. Replace line 3 in the signing algorithm of Fig. 1 by the conditional preimage sampling, i.e., $d \leftarrow \mathsf{SampleDom}_\rho(1^\lambda)$ conditioned on $f_\rho(d) = h$. In the anonymity game, observe that the signature is independent of b, so the adversary can only output the correct b with probability $1/2$. It remains to consider how to bound the difference between the original and the modified signing oracle.

To get Eq. (5), we use the statistical distance $\epsilon^{pre}_{\mathsf{RPSF}}$. For both $b = 0$ or $b = 1$, the statistical distance for each query is $\epsilon^{pre}_{\mathsf{RPSF}}$, and with q_c, the overall difference in winning in the unmodified and the modified game is at most $q_c \epsilon^{pre}_{\mathsf{RPSF}}$.

For Eq. (6), we use the Kullback-Leibler divergence. For both $b = 0$ and $b = 1$, the only change is in the signing oracle, and the difference between signing oracle and the domain sampling is given by the Kullback-Leibler divergence with $\epsilon^{KL\text{-}pre}_{\mathsf{RPSF}}$. By application of Lemma 3 for q_c queries, we get the desired relation. □

6 AOS Ring Signatures

Abe et al. [1] provide simple ring signatures from any Σ-protocol using a circular version of the Fiat-Shamir transform [32], the AOS-transform. It is used in several linear-size ring signatures [13,45]. Its security was recently proven by Yuen et al. [63] for canonical verification and established for general AOS-like ring signatures by Borin et al. [13]. However, both proofs are in the ROM only.

The specification of AOS ring signatures is in Fig. 2. Given a Σ-protocol, we refer to the ring signature obtained using the AOS-framework as $\mathsf{AOS}(\Sigma)$. We will show SUF-CRA security of AOS ring signatures in the QROM to complement our results on the RPSF-based RSS. The techniques we use are unrelated to the ones developed in Section 3. First, we use Lemma 4 to relate the SUF-CRA security to the UF-NRA security. In Section 6.2, we employ the measure-and-reprogram technique from Theorem 1 to generically relate the UF-NRA security to impersonation security of the Σ-protocol. For *C&O* Σ-protocols, we give a tighter alternative in Section 6.3, generalizing the approach from [28]. In this section, we assume perfect correctness for simplicity, but the technique also works for imperfect correctness.

$\mathsf{Sign}(\mathsf{sk} = w, \rho = \{\mathsf{inst}_1, \ldots, \mathsf{inst}_N\}, m)$

1 : // let $i \in [N]$ be the index s.t. $\mu(w) = \mathsf{inst}_i$
2 : $(\mathsf{com}_i, \mathsf{state}_i) \leftarrow \Sigma.\mathsf{P}_1(\mathsf{inst}_i)$
3 : **for** $j = i+1, \ldots, N, 1, \ldots, i-1$
4 : $\quad \mathsf{ch}_j \leftarrow \mathsf{H}(j, \rho, \mathsf{com}_{j-1}, m)$
5 : $\quad (\mathsf{com}_j, \mathsf{rsp}_j) \leftarrow \mathsf{Sim}(\mathsf{inst}_j, \mathsf{ch}_j)$
6 : $\mathsf{ch}_i \leftarrow \mathsf{H}(i, \rho, \mathsf{com}_{i-1}, m)$
7 : $\mathsf{rsp} \leftarrow \Sigma.\mathsf{P}_2(\mathsf{state}, w, \mathsf{ch}_i)$
8 : $\sigma \leftarrow ((\mathsf{com}_1, \mathsf{rsp}_1), \ldots, (\mathsf{com}_N, \mathsf{rsp}_N))$
9 : **return** σ

$\mathsf{Vf}(\rho = \{\mathsf{inst}_1, \ldots, \mathsf{inst}_N\}, m, \sigma)$

1 : $\mathsf{ch}_1 = \mathsf{H}(1, \rho, \mathsf{com}_N, m)$
2 : **for** $j = 1, \ldots, N$
3 : $\quad$ **if** $\neg \mathsf{ID}.\mathsf{V}_2(\mathsf{inst}_j, \mathsf{com}_j, \mathsf{ch}_j, \mathsf{rsp}_j)$
4 : $\quad\quad$ **return** false
5 : $\quad \mathsf{ch}_{j+1} \leftarrow \mathsf{H}(j+1, \rho, \mathsf{com}_j, m)$
6 : **return** true

Fig. 2. Construction of a ring signature $\mathsf{AOS}(\Sigma)$ from a Σ-protocol Σ and its HVZK simulator Sim using [1]. The construction can be modified if Σ has commitment recoverability. In the modified version, ch_1 can be sent instead of all the commitments, and the commitments and challenges are computed on the fly. Finally, there would be a consistency check, if ch_1 is the challenge that can be recovered from the last commitment. Stp only fixes the maximal ring size and KGen is identical to $\Sigma.\mathsf{Gen}$.

6.1 SUF-CRA to UF-NRA

Theorem 8. *Let Σ be a special-sound Σ-protocol with HVZK simulator Sim and simulator commitment min-entropy*[6] *$\beta(\lambda)$. Consider the AOS RSS with maximal ring size $\kappa = \kappa(\lambda) \geq N$. For any* SUF-CRA *adversary $\mathcal{A}$ making at most q_s signing queries and* ANON *adversary $\mathcal{B}$ making at most q_c challenge queries, each making at most q_{H} quantum queries to the RO H, there exists a* UF-NRA *adversary $\mathcal{A}_{\text{UF-NRA}}$, a* CUR *adversary $\mathcal{A}_{\text{CUR}}$ and a* W-REC *adversary $\mathcal{A}_{\text{W-REC}}$ such that $\mathsf{Adv}^{\text{SUF-CRA}}_{\mathcal{A},\mathsf{AOS}(\Sigma),N,q_s}(\lambda)$ and $\mathsf{Adv}^{\text{ANON}}_{\mathcal{B},\mathsf{AOS}(\Sigma),N,q_c}(\lambda)$ are bounded as in Table 5 and the parameters are defined as in Table 6.*

Proof (Sketch). The full proof can be found in Appendix B.4. For unforgeability we use Lemma 4 to reprogram every signature query with uniform challenges, which can be done assuming sufficient min-entropy in the hash input. This introduces the typical adaptive reprogramming error (ϵ_{repr}), and an additive error term (ϵ'_{repr}) to check whether the final forgery is valid for the non-reprogrammed random oracle. Then, we use the HVZK simulator to remove use of the secret key which can be done using Rényi divergence and statistical distance.

For anonymity we reprogram the random oracle, and then we use the HVZK simulator. Statistical HVZK gives a straightforward bound based on the number of challenge queries. Alternatively, the Kullback-Leibler divergence can be used, which gives a direct bound using Lemma 3. □

[6] If the min-entropy is too small, we can either add commitment entropy by appending a random string, or we can instantiate the AOS construction with a *salted* RO, resulting in a higher effective commitment min-entropy.

Table 5. Unforgeability and anonymity bounds in Theorem 8 using statistical distance (SD), Rényi divergence (RD) or Kullback-Leibler divergence (KL) for $\alpha \in (1, \infty]$.

Security	Method	Bound
$\mathsf{Adv}^{\text{SUF-CRA}}_{\mathcal{A},\mathsf{AOS}(\Sigma),N,q_s}(\lambda)$	SD	$\epsilon_{\text{UF-NRA}} + q_s \cdot \epsilon_{\Sigma}^{\text{HVZK}} + \epsilon_{repr} + \epsilon'_{repr}$
	RD	$\left[(\delta_{\Sigma}^{\alpha\text{-HVZK}})^{q_s} \epsilon_{\text{UF-NRA}}\right]^{\frac{\alpha-1}{\alpha}} + \epsilon_{repr} + \epsilon'_{repr}$
$\mathsf{Adv}^{\text{ANON}}_{\mathcal{B},\mathsf{AOS}(\Sigma),N,q_c}(\lambda)$	SD	$q_c \cdot \epsilon_{\Sigma}^{\text{HVZK}} + \epsilon''_{repr}$
	KL	$\sqrt{(q_c \cdot \epsilon_{\Sigma}^{KL\text{-HVZK}})/2} + \epsilon''_{repr}$

Table 6. Parameters from Theorem 8 where $\epsilon_{\text{UF-NRA}} := \mathsf{Adv}^{\text{UF-NRA}}_{\mathcal{A}_{\text{UF-NRA}},\mathsf{AOS}(\Sigma),N}(\lambda)$.

	Definition
ϵ_{repr}	$\frac{3q_s}{2}\sqrt{(q_{\mathsf{H}} + \kappa \cdot q_s + N) \cdot 2^{-\beta(\lambda)}}$
ϵ'_{repr}	$N \cdot \left(\mathsf{Adv}^{\text{CUR}}_{\mathcal{A}_{\text{CUR}},\Sigma}(\lambda) + \mathsf{Adv}^{\text{W-REC}}_{\mathcal{A}_{\text{W-REC}},\Sigma}(\lambda)\right) + 20(q_{\mathsf{H}} + N \cdot q_s + N)^3/\lvert\mathcal{C}\rvert$
ϵ''_{repr}	$\frac{3q_c}{2}\sqrt{(q_{\mathsf{H}} + \kappa \cdot q_c) \cdot 2^{-\beta(\lambda)}}$

Proof in the ROM As in Section 5.1, the main changes arise from natural transitioning between collision resistance in the QROM to the ROM and the adaptive reprogramming. Both are only additive terms, and when viewing in the ROM, smaller choices for the min-entropy and the challenge space are likely possible. A proof of UF-CRA in the ROM can be found in [13, Proposition A.3].

6.2 UF-NRA: General Strategy using Measure-and-Reprogram

When considering the classical reduction from SUF-CRA to UF-NRA, a time-ordered list of the queries is essential. The ROM naturally provides this time-ordered list (and only has a guessing cost of $q_{\mathsf{H}} \cdot (q_{\mathsf{H}} - 1)$ [13, In the proof of Proposition A.3]), but unfortunately not the QROM. We can get such a list using the measure-and-reprogram technique [26]. We can get a query order for a subset of queries, allowing us to reprogram these queries in that specific order.

Theorem 9 (UF-NRA using Measure-and-Reprogram 2.0). *Let Σ be a Σ-protocol. For any* UF-NRA *adversary $\mathcal{A}$ making at most q_{H} quantum queries to H, there exists an impersonation adversary $\mathcal{B}$ that can convince an honest verifier V such that*

$$\mathsf{Adv}^{\text{UF-NRA}}_{\mathcal{A},\mathsf{AOS}(\Sigma),N}(\lambda) \leq N \cdot (2(q_{\mathsf{H}} + N) + 1)^{2N} \cdot \mathsf{Adv}^{\text{IMP}}_{\mathcal{B},\Sigma}(\lambda).$$

Proof (Sketch). The full proof can be found in Appendix B.5. We use Theorem 1 for an ordering of all N queries used within the verification of the adversaries final forgery. There always exists an index in this list, where we can successfully

include an interactive impersonation adversary against Σ. The first message must be made before the respective challenge for it is fixed by the previous ring member, and with the list, such an index always exists. The application of Theorem 1 yields the multiplicative loss in the order of q_{H}^{2N}. □

6.3 UF-NRA: Commit-and-Open Sigma Protocols

In this section, we consider commit-and-open ($C\&O$) Σ-protocols to construct AOS RSS. In a $C\&O$ Σ-protocol, the prover commits to responses within the first message of the protocol. The challenge defines which responses need to be opened in the third protocol message. Formally, we define them as follows.

Definition 2 (Commit-and-Open Sigma Protocols). *A commit-and-open Σ-protocol Σ is a Σ-protocol with a special form that uses a hash function H. In this protocol, the first message consists of ℓ commitments $y_1, \ldots, y_\ell$ computed as $y_i = \mathsf{H}(m_i)$ for $m_i \in \mathcal{M}$ and possibly an additional string $a_\circ$.[7] The challenge is then picked uniformly at random from $\mathcal{C} \subseteq 2^{[\ell]}$. This identifies the responses as $\boldsymbol{m}_c = (m_i)_{i \in c}$. The verifier accepts iff $\mathsf{H}(m_i) = y_i$ for all $i \in c$ and some given predicate $\mathsf{V}(\mathsf{inst}, c, \boldsymbol{m}_c, a_\circ)$ is satisfied.*

Don et al. [29] showed that $C\&O$ Σ-protocols have online extractability after application of the Fiat-Shamir transform. We introduce their soundness definition for Σ-protocols, which will also be essential in our proof. We recall the definition of $\mathfrak{S}^*$-soundness for $C\&O$ Σ-protocols from [28].

Consider $\mathcal{C} \subseteq 2^{[\ell]}$, and let $\mathfrak{S} \subseteq 2^{\mathcal{C}}$ be an arbitrary non-empty, monotone increasing (i.e. $S \in \mathfrak{S} \wedge S \subseteq S' \Rightarrow S' \in \mathfrak{S}$) set of subsets $S \subseteq \mathcal{C}$. Let $\mathfrak{S}_{\min} := \{S \in \mathfrak{S} \mid S' \subsetneq S \Rightarrow S' \notin \mathfrak{S}\}$. As an example, consider $\mathfrak{S}$ to be $\{S \subseteq \mathcal{C} \mid |S| \geq k\}$, then $\mathfrak{S}_{min}$ is defined as the subsets of $\mathcal{C}$ with exactly k elements.

Definition 3 ([29, Def. 5.2]). *A $C\&O$ Σ-protocol Σ is $\mathfrak{S}$-sound* if there exists an efficient deterministic algorithm $\mathsf{Ext}^*_{\mathfrak{S}}(\mathsf{inst}, m_1, \ldots, m_\ell, a_\circ)$ that takes on input an instance $\mathsf{inst} \in \mathcal{I}$, messages $m_1, \ldots, m_\ell \in \mathcal{M} \cup \{\bot\}$ and a string $a_\circ$, outputs a witness for inst, if $\exists S \in \mathfrak{S}$ such that $\mathsf{V}(\mathsf{inst}, c, \boldsymbol{m}_c, a_\circ)$ for all $c \in S$.*

Informally, this soundness implies the existence of an extractor capable of recovering a witness if the transcript is valid for many challenges. This is modeled using $\mathfrak{S}_{\min}$ *without* explicitly being given the set S. Similar to [29], we define

$$sz^{\mathfrak{S}}_{triv} := \max_{\hat{S} \notin \mathfrak{S}} |\hat{S}|.$$

capturing the number of challenges a prover can successfully answer by first picking a set $\hat{S} \notin \mathfrak{S}$ of challenges, and then preparing a message $\hat{\boldsymbol{m}}$ and a_0 such that $\mathsf{V}(\mathsf{inst}, c, \hat{\boldsymbol{m}}_c, a_0)$ holds if $c \in \hat{S}$. If an adversary prepares the same for an $S \in \mathfrak{S}$, running the soundness-extractor retrieves the witness. [29] defines the probability of a successful attack by dividing $sz^{\mathfrak{S}}_{triv}$ by $|\mathcal{C}|$ as $p^{\mathfrak{S}}_{triv}$.

[7] Note that $m_i \in \mathcal{M}$ may contain some additional randomness.

UF-NRA using *C&O* Σ-Protocols Our goal is to show that AOS RSSs are UF-NRA secure by utilizing the online-extractability of *C&O* Σ-protocols. The main tool for this analysis uses the framework introduced in Section 2.3. We follow a similar procedure to that of Don et al. [28, Section 4] when proving the online extractability of *C&O* Σ-protocols under the Fiat-Shamir transformation. (Alternatively, the formalism of [19] could be used.) First, we introduce three database properties, similar to those in [28]. The properties capture collisions, limited size, and the event where a proof can be produced, and no extraction is possible. The size property can be defined as

$$\mathsf{SZ}_{\leq s} := \{D \,|\, |\{x|D(x) \neq \bot\}| \leq s\}\,.$$

We now define the prover-success-extractor-fail property. Consider a set of N instances $\{\mathsf{inst}_1, \ldots, \mathsf{inst}_N\} \subseteq \mathcal{I}$.We define the sets $\mathcal{M}$ for messages, $\mathcal{C}$ for challenges, and $\mathcal{Y}$ for commitments referring to the sets from the *C&O* Σ-protocols. We define commitment queries as $D(m)$ for $m \in \mathcal{M}$, returning a commitment $y \in \mathcal{Y}$ and extend this notation to vectors of messages $\mathbf{m}$. For a database D and a commitment $y \in \mathcal{Y}$, we define $D^{-1}(y)$ to be the smallest $m \in \mathcal{M}$ with $D(m) = y$. If no such m exists, we define $D^{-1}(y) = \bot$, as well as $D^{-1}(\bot) = \bot$. We denote challenge queries as $D(\beta, \rho, \boldsymbol{y}, \mu)$ where $\boldsymbol{y} \in \mathcal{Y}^\ell$, μ acts as auxiliary input, and ρ is a collection of instances. The output is a string from $\mathcal{Y}$, which means that we need an additional map $\gamma : \mathcal{Y} \to \mathcal{C}$. With this, we can define the remaining two database properties of collision[8]

$$\mathsf{CL} := \left\{D \,\middle|\, \exists x \neq x' \in \mathcal{X} : D(x) \neq \bot \neq D(x') \wedge \gamma \circ D(x) = \gamma \circ D(x')\right\}$$

and adversarial success ($\rho' \subseteq \rho \subseteq \mathrm{Supp}\big(\Sigma.\mathsf{Gen}(1^\lambda)\big)$, $N' = |\rho'| \leq |\rho| = N$)

$$\mathsf{SUC} := \left\{ D \,\middle|\, \begin{array}{c} \exists \mu \text{ and } \boldsymbol{y}_\beta \in \mathcal{Y}^\ell \text{ s.t. } \forall \beta \in [N'] : \\ \mathsf{V}(\mathsf{inst}_\beta, c_\beta, \boldsymbol{m}_{\beta,c_\beta}) \wedge (\mathsf{inst}_\beta, \mathsf{Ext}^*(\mathsf{inst}_\beta, \boldsymbol{m}_\beta)) \notin R \\ \text{where } c_{\beta+1} := \gamma \circ D(\beta+1, \rho', \boldsymbol{y}_\beta, \mu) \wedge \boldsymbol{m}_\beta := D^{-1}(\boldsymbol{y}_\beta) \end{array} \right\}.$$

Our intermediate goal is to bound the transition capacity from an empty compressed database to satisfy either SUC or CL after q queries.

Lemma 7. *Let* $\gamma_{cl} := \max_{c \in \mathcal{C}} |\{y \in \mathcal{Y} \mid \gamma(y) = c\}|$ *and* $N < q$*, then*

$$[\![\bot \Rightarrow^q \mathsf{SUC} \cup \mathsf{CL}]\!] \leq q\sqrt{10/|\mathcal{Y}|}\left(\sqrt{(q-1)\gamma_{cl}} + \sqrt{\max\left\{(q-1)\ell, \gamma_{cl}(sz^{\mathfrak{S}}_{triv})^N\right\}}\right).$$

Proof. Using the same arguments as in [28, Proof of Lemma 4], we get that

$$[\![\bot \Rightarrow^q \mathsf{SUC} \cup \mathsf{CL}]\!] \leq \sum_{s=0}^{q-1}\Big([\![\mathsf{SZ}_{\leq s} \setminus \mathsf{CL} \to \mathsf{CL}]\!] + [\![\mathsf{SZ}_{\leq s} \setminus \mathsf{SUC} \setminus \mathsf{CL} \to \mathsf{SUC}]\!]\Big).$$

[8] This collision property captures slightly more than what we need, as we could also separate between collisions for commitments and challenges.

Compared to [28], we have implicitly removed the first term, as it vanishes regardless, and we have not simplified the last term, as we explicitly need the collision property. The first term is bounded similarly to [20, Example 5.28], but we use Lemma 8 to get a tighter bound. The second term is bounded by Lemma 10. Finally, we combine the results and use that $s \leq q - 1$. □

Lemma 8. *Let* $\gamma_{cl} := \max_{c\in\mathcal{C}} |\{y \in \mathcal{Y} \mid \gamma(y) = c\}|$*, then*

$$[\![\mathsf{SZ}_{\leq s} \setminus \mathsf{CL} \to \mathsf{CL}]\!] \leq \sqrt{10}\sqrt{s \cdot \gamma_{cl}/|\mathcal{Y}|}.$$

Similar results can be found in [20, Example 5.28], so we defer the proof to Appendix B.6. We can now finalize Theorem 8 using the previously established bound on the quantum transition capacity.

Lemma 9. *Let* $\gamma_{cl} := \max_{c\in\mathcal{C}} |\{y \in \mathcal{Y} \mid \gamma(y) = c\}|$*, then for any oracle quantum algorithm* $\mathcal{A}$ *with query complexity* q*, it holds that*

$$\Pr_{\mathcal{A}}[D \in \mathsf{CL}] \leq 10q^2(q-1) \cdot \gamma_{cl}/|\mathcal{Y}|.$$

Proof. Consider $[\![\bot \Rightarrow^q \mathsf{CL}]\!] \leq \sum_{s=0}^{q-1} [\![\mathsf{SZ}_{\leq s} \setminus \mathsf{CL} \to \mathsf{CL}]\!]$. Using Lemma 8 we get

$$\sum_{s=0}^{q-1} [\![\mathsf{SZ}_{\leq s} \setminus \mathsf{CL} \to \mathsf{CL}]\!] \leq q\sqrt{10}\sqrt{(q-1) \cdot \gamma_{cl}/|\mathcal{Y}|}.$$

Squaring this term yields the upper bound for any $\mathcal{A}$ with q queries. □

Lemma 10. *Let* $\gamma_{cl} := \max_{c\in\mathcal{C}} |\{y \in \mathcal{Y} \mid \gamma(y) = c\}|$*, then*

$$[\![\mathsf{SZ}_{\leq s} \setminus \mathsf{SUC} \setminus \mathsf{CL} \to \mathsf{SUC}]\!] \leq \sqrt{10}\sqrt{\max\left\{N \cdot \ell, s \cdot \ell, \gamma_{cl} \cdot (sz^{\mathfrak{S}}_{triv})^N\right\}/|\mathcal{Y}|}.$$

A similar result can be found in [28, Proof of Lemma 4.1], so we defer the full proof to the Appendix B.7. The main modifications are based on having to account for more potential targets due to the ring size N. Specifically, $sz^{\mathfrak{S}}_{triv}$ captures the number of challenges that can be answered without allowing for extraction. This size is now propagated through the entire ring, giving a total number of potentially valid targets of $(sz^{\mathfrak{S}}_{triv})^N$, where it was only $sz^{\mathfrak{S}}_{triv}$ in [28].

Theorem 10. *Let* Σ *be a* $\mathfrak{S}^*$*-sound C&O* Σ*-protocol with challenge space* $\mathcal{C} = \mathcal{C}_\lambda$*,* $\ell = \ell(1^\lambda)$ *commitments,* $\mathsf{H} : \{0,1\}^* \to \mathcal{Y}$ *be a hash function modeled in the QROM and* $\gamma : \mathcal{Y}^\ell \to \mathcal{C}_\lambda$ *be a function with* $\gamma_{cl} := \max_{c\in\mathcal{C}} |\{y \in \mathcal{Y} \mid \gamma(y) = c\}|$*. Set* $\omega = \omega(\lambda) := \max_{c\in\mathcal{C}_\lambda} |c|$*. Furthermore, let* $N \in \mathbb{N}$ *be fixed with* $N < q_{\mathsf{H}}$*. For any* UF-NRA *adversary* $\mathcal{A}$ *making at most* q_{H} *quantum queries to* H*, there exists an efficient adversary* $\mathcal{B}$ *against witness-recovery of the* Σ *with*

$$\begin{aligned}
\mathsf{Adv}^{\text{UF-NRA}}_{\mathcal{A},\mathsf{AOS}(\Sigma),N}(\lambda) &\leq N \cdot \mathsf{Adv}^{\text{W-REC}}_{\mathcal{B},\Sigma}(\lambda) + 2 \cdot N(\omega(\lambda)+1) \cdot |\mathcal{Y}|^{-1} \\
&\quad + q^2_{\mathsf{H}} 10/|\mathcal{Y}| \left(\sqrt{(q_{\mathsf{H}}-1) \cdot \gamma_{cl}} + \sqrt{\max\left\{(q_{\mathsf{H}}-1) \cdot \ell, \gamma_{cl} \cdot (sz^{\mathfrak{S}}_{triv})^N\right\}}\right)^2. \\
&\leq N \cdot \mathsf{Adv}^{\text{W-REC}}_{\mathcal{B},\Sigma}(\lambda) + 2 \cdot N(\omega(\lambda)+1) \cdot |\mathcal{Y}|^{-1} \\
&\quad + q^2_{\mathsf{H}}(q_{\mathsf{H}}-1)\ell 10\gamma_{cl}/|\mathcal{Y}| \left(1 + 2(sz^{\mathfrak{S}}_{triv})^{N/2} + (sz^{\mathfrak{S}}_{triv})^N\right).
\end{aligned}$$

Proof. Let $\mathcal{A}$ be an UF-NRA-adversary. We will construct a witness-recovery adversary $\mathcal{B}$ for Σ that succeeds if $\mathcal{A}$ produces a forgery in the UF-NRA game. $\mathcal{B}$ is given inst where $(\mathsf{inst}, w) \leftarrow \mathsf{Gen}(1^\lambda)$. Now, $\mathcal{B}$ guesses an index $i^* \leftarrow [N]$. For all other $i \in [N] \backslash \{i^*\}$, $\mathcal{B}$ computes $(\mathsf{inst}_i, w_i) \leftarrow \mathsf{Gen}(1^\lambda)$. Set $\rho = \{\mathsf{inst}_1, \ldots, \mathsf{inst}_N\}$. For simplicity, we assume that all challenges $c \in \mathcal{C}$ are of the same length.

$\mathcal{B}$ gives ρ to $\mathcal{A}$ and plays the role of the challenger simulating the random oracle of $\mathcal{A}$ as a compressed oracle. We denote this by $\mathcal{A}^{\mathsf{CO}}$. After making q_{H} queries to the random oracle H, $\mathcal{A}$ outputs (ρ^*, m^*, σ^*) where $\rho^* \subseteq \rho$. If i^* is not part of the ring, we abort. Now, each σ_i with $i \in [N^*]$ where $\sigma^* = (\sigma_1, \ldots, \sigma_{N^*})$ consists of commitments $\boldsymbol{y}_i \in \mathcal{Y}^\ell$ and messages $\boldsymbol{m}_{i,c_i} = (m_{i,j})_{j \in c_i}$. Now, we obtain $c_{i+1} = \gamma \circ \mathsf{H}(i+1, \rho^*, \boldsymbol{y}_i, \mu)$ as well as $\boldsymbol{h}_{i,c_i} = \mathsf{H}(\mathbf{m}_{i,c_i})$. For a signature to be valid, the relation $\tilde{R}_{\rho^*}$ must hold for the triple $(\sigma^*, \boldsymbol{c}, \boldsymbol{h})$ which is defined as

$$\forall i \in [N^*] : \boldsymbol{h}_{i,c_i} = \boldsymbol{y}_{i,c_i} \quad \wedge \quad \mathsf{V}(\rho_i, c_i, \boldsymbol{m}_{i,c_i}).$$

Next, we measure the internal state of the compressed oracle to obtain a database D. Let j^* be the index with which i^* is identified in ρ^*. Now, we obtain $m_{j,c} = D^{-1}(y_{j,c})$ for $c \in \mathcal{C}$ and run $\mathsf{Ext}^*(\rho^*, \boldsymbol{m}_{j^*})$ to obtain a witness w.

Defining $\hat{\boldsymbol{c}}$ and $\hat{\boldsymbol{h}}$ like $\boldsymbol{c}$ and $\boldsymbol{h}$, but using D in place of H, we get

$$\begin{aligned}\mathsf{Adv}^{\text{UF-NRA}}_{\mathcal{A},\mathsf{RSig},N}(\lambda) &:= \Pr[\mathsf{Vf}(\rho^*, m^*, \sigma^*) \wedge \rho^* \subseteq \rho] \leq \Pr\Big[(\sigma^*, \boldsymbol{c}, \boldsymbol{h}) \in \tilde{R}_{\rho^*} \wedge \rho^* \subseteq \rho\Big] \\ &\leq \Pr\Big[(\sigma^*, \hat{\boldsymbol{c}}, \hat{\boldsymbol{h}}) \in \tilde{R}_{\rho^*} \wedge \rho^* \subseteq \rho\Big] + 2 \cdot N^*(\omega(\lambda) + 1) \cdot |\mathcal{Y}|^{-1}\end{aligned}$$

In the last step, we have used Lemma 5. Now, we bound the first term by distinguishing whether the database D satisfies $\mathsf{SUC} \cup \mathsf{CL}$. For simplicity, we drop the explicit notation for how (ρ^*, m^*, σ^*) are generated. We bound

$$\begin{aligned}\Pr\Big[(\sigma^*, \hat{\boldsymbol{c}}, \hat{\boldsymbol{h}}) \in \tilde{R}_{\rho^*} \wedge \rho^* \subseteq \rho\Big] &\leq \Pr\Big[(\sigma^*, \hat{\boldsymbol{c}}, \hat{\boldsymbol{h}}) \in \tilde{R}_{\rho^*} \wedge \rho^* \subseteq \rho \wedge D \in \mathsf{SUC} \cup \mathsf{CL}\Big] \\ &+ \Pr\Big[(\sigma^*, \hat{\boldsymbol{c}}, \hat{\boldsymbol{h}}) \in \tilde{R}_{\rho^*} \wedge \rho^* \subseteq \rho \wedge D \notin \mathsf{SUC} \cup \mathsf{CL}\Big].\end{aligned}$$

For the first case, we have

$$\begin{aligned}&\Pr\Big[(\sigma^*, \hat{\boldsymbol{c}}, \hat{\boldsymbol{h}}) \in \tilde{R}_{\rho^*} \wedge \rho^* \subseteq \rho \wedge D \in \mathsf{SUC} \cup \mathsf{CL}\Big] \leq \Pr[D \in \mathsf{SUC} \cup \mathsf{CL}] \\ &\leq [\![\bot \Rightarrow^{q_{\mathsf{H}}} \mathsf{SUC} \cup \mathsf{CL}]\!]^2.\end{aligned}$$

In the second case, we get that $(\sigma^*, \hat{\boldsymbol{c}}, \hat{\boldsymbol{h}}) \in \tilde{R}_{\rho^*}$ and $D \notin \mathsf{SUC} \cup \mathsf{CL} \subseteq \mathsf{SUC}$. Considering the definition of SUC, there exists $i \in [N^*] : (\mathsf{inst}_i, \mathsf{Ext}^*(\mathsf{inst}_i, D^{-1}(\boldsymbol{y}_i))) \in R$ succeeds because the other terms are all satisfied due to $\tilde{R}_{\rho^*}$. Formally,

$$\begin{aligned}&\Pr\Big[(\sigma^*, \hat{\boldsymbol{c}}, \hat{\boldsymbol{h}}) \in \tilde{R}_{\rho^*} \wedge \rho^* \subseteq \rho \wedge D \notin \mathsf{SUC} \cup \mathsf{CL}\Big] \\ &\quad \leq \Pr\big[\exists i \in [N^*] : (\mathsf{inst}_i, \mathsf{Ext}^*(\mathsf{inst}_i, D^{-1}(\boldsymbol{y}_i))) \in R\big].\end{aligned}$$

With probability $\frac{1}{N^*}$, this is the index where the challenge instance of $\mathcal{B}$ is placed, and $\mathcal{B}$ outputs $\mathsf{Ext}^*(\mathsf{inst}_i, \boldsymbol{m}_i)$ as a witness. As the index where the challenge

instance was placed is used in the signature with probability $\frac{N^*}{N}$, we get

$$N^{-1}\Pr\left[(\sigma^*, \hat{\boldsymbol{c}}, \hat{\boldsymbol{h}}) \text{ satisfies } \tilde{R}_{\rho^*} \wedge D \notin \mathsf{SUC} \cup \mathsf{CL}\right] \leq \mathsf{Adv}^{\text{W-REC}}_{\mathcal{B},\Sigma}(\lambda),$$

and combining these results with Lemma 7 shows our final bound. □

In [28], the results for $C\&O$ protocols are generalized to allow for Merkle tree based vector commitments. We perform the same generalization for the AOS transform, see Appendix D.

7 Instantiation and Discussion of Ring Signatures

AOS-based Ring Signature. There are several AOS-based RSS in the literature. Erebor [13] and MayoRS [45] are recent examples. Neither is based on a $C\&O$ Σ-protocol, but Theorem 9 yields the first QROM security bound for both. The bounds are not tight, but at least rule out fundamental quantum vulnerabilities.

Explicit examples for $C\&O$ Σ-protocols are ZKBOO [37] and ZKB++ [16] which are the underlying Σ-protocols in Fish and Picnic [16]. More generally, $C\&O$ Σ-protocols can be constructed using the MPC-in-the-head paradigm [44].

RPSF-based Ring Signature. As an example, consider the ring signature Gandalf [35]. The translation of Gandalf-like ring signatures into our RPSF framework is relatively straightforward (for details see Appendix E.2). Compared to the classical proof technique, we get the modified hash-collision term from the QROM, the additional additive adaptive reprogramming term that requires sufficiently many salt bits k and the typical square-root loss when reducing to the one-wayness via measure-and-reprogram. Our framework can also be applied (both for Theorem 6 and Theorem 7) to [62], a lattice-based ring signature directly based on [36].

Comparison and Trade-offs. The generic AOS bound is exponential in ring size N and is not expected to be tight, but gives meaningful guarantees for small rings. The $C\&O$ bound is multiplicatively tight and avoids exponential scaling when $sz_{\text{triv}} = 1$, e.g., for protocols with 2-special soundness. RPSF-based bounds are tight but constructing RPSFs for large rings is non-trivial, as the assumed hardness of finding collisions is impacted by the ring size.

Computational Assumptions. In our work, we relied on statistical HVZK properties of Σ-protocols. For Σ-protocols in AOS, it is sufficient to consider weak special HVZK properties for unforgeability, i.e., an adversary given a transcript from the simulator or the actual algorithm cannot distinguish them. Analogously, for RPSFs, computational indistinguishability between presampling and conditional domain sampling would suffice. For domain sampling, computational properties suffice when considering Theorem 6, but not when considering Theorem 7.

Anonymity from computational assumptions requires additional access to *the witness/trapdoor* (called strong HVZK of Σ-protocols [13]).

Acknowledgments. We thank the anonymous reviewers of Eurocrypt 2026, Crypto 2026 and Asiacrypt 2026 for their valuable feedback. We used Claude for language editing and to improve presentation; all technical content and proofs are the authors' own. The authors acknowledge support from the Danish Ministry of Defense Acquisition and Logistics Organization (FMI). CM acknowledges support by the Independent Research Fund Denmark via a DFF Sapere Aude grant (IM-3PQC, grant ID 10.46540/2064-00034B).

References

1. Abe, M., Ohkubo, M., Suzuki, K.: 1-out-of-n signatures from a variety of keys. In: Zheng, Y. (ed.) ASIACRYPT 2002. LNCS, vol. 2501, pp. 415–432. Springer, Berlin, Heidelberg (Dec 2002). https://doi.org/10.1007/3-540-36178-2_26
2. Alagic, G., Bai, C., Katz, J., Majenz, C.: Post-quantum security of the Even-Mansour cipher. In: Dunkelman, O., Dziembowski, S. (eds.) EUROCRYPT 2022, Part III. LNCS, vol. 13277, pp. 458–487. Springer, Cham (May / Jun 2022). https://doi.org/10.1007/978-3-031-07082-2_17
3. Bai, S., Langlois, A., Lepoint, T., Stehlé, D., Steinfeld, R.: Improved security proofs in lattice-based cryptography: Using the Rényi divergence rather than the statistical distance. In: Iwata, T., Cheon, J.H. (eds.) ASIACRYPT 2015, Part I. LNCS, vol. 9452, pp. 3–24. Springer, Berlin, Heidelberg (Nov / Dec 2015). https://doi.org/10.1007/978-3-662-48797-6_1
4. Banaszczyk, W.: New bounds in some transference theorems in the geometry of numbers. Mathematische Annalen **296**(1), 625–635 (1993)
5. Barbosa, M., Barthe, G., Doczkal, C., Don, J., Fehr, S., Grégoire, B., Huang, Y.H., Hülsing, A., Lee, Y., Wu, X.: Fixing and mechanizing the security proof of Fiat-Shamir with aborts and Dilithium. In: Handschuh, H., Lysyanskaya, A. (eds.) CRYPTO 2023, Part V. LNCS, vol. 14085, pp. 358–389. Springer, Cham (Aug 2023). https://doi.org/10.1007/978-3-031-38554-4_12
6. Barnes, R., Bhargavan, K., Lipp, B., Wood, C.A.: Hybrid Public Key Encryption. RFC 9180 (Feb 2022). https://doi.org/10.17487/RFC9180, https://www.rfc-editor.org/info/rfc9180
7. Belovs, A.: Quantum algorithms for classical probability distributions. arXiv preprint arXiv:1904.02192 (2019)
8. Bender, A., Katz, J., Morselli, R.: Ring signatures: Stronger definitions, and constructions without random oracles. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 60–79. Springer, Berlin, Heidelberg (Mar 2006). https://doi.org/10.1007/11681878_4
9. Bender, A., Katz, J., Morselli, R.: Ring signatures: Stronger definitions, and constructions without random oracles. Journal of Cryptology **22**(1), 114–138 (Jan 2009). https://doi.org/10.1007/s00145-007-9011-9
10. Bernstein, D.J.: Multi-ciphertext security degradation for lattices. Cryptology ePrint Archive, Report 2022/1580 (2022), https://eprint.iacr.org/2022/1580
11. Beullens, W., Katsumata, S., Pintore, F.: Calamari and Falafl: Logarithmic (linkable) ring signatures from isogenies and lattices. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020, Part II. LNCS, vol. 12492, pp. 464–492. Springer, Cham (Dec 2020). https://doi.org/10.1007/978-3-030-64834-3_16
12. Boneh, D., Dagdelen, Ö., Fischlin, M., Lehmann, A., Schaffner, C., Zhandry, M.: Random oracles in a quantum world. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 41–69. Springer, Berlin, Heidelberg (Dec 2011). https://doi.org/10.1007/978-3-642-25385-0_3
13. Borin, G., Lai, Y.F., Leroux, A.: Erebor and durian: Full anonymous ring signatures from quaternions and isogenies. CiC **1**(4), 4 (2024). https://doi.org/10.62056/ava3zivrzn
14. Brakerski, Z., Kalai, Y.T.: A framework for efficient signatures, ring signatures and identity based encryption in the standard model. Cryptology ePrint Archive, Report 2010/086 (2010), https://eprint.iacr.org/2010/086

15. Brendel, J., Fiedler, R., Günther, F., Janson, C., Stebila, D.: Post-quantum asynchronous deniable key exchange and the Signal handshake. In: Hanaoka, G., Shikata, J., Watanabe, Y. (eds.) PKC 2022, Part II. LNCS, vol. 13178, pp. 3–34. Springer, Cham (Mar 2022). https://doi.org/10.1007/978-3-030-97131-1_1
16. Chase, M., Derler, D., Goldfeder, S., Orlandi, C., Ramacher, S., Rechberger, C., Slamanig, D., Zaverucha, G.: Post-quantum zero-knowledge and signatures from symmetric-key primitives. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) ACM CCS 2017. pp. 1825–1842. ACM Press (Oct / Nov 2017). https://doi.org/10.1145/3133956.3133997
17. Chatterjee, R., Chung, K.M., Liang, X., Malavolta, G.: A note on the post-quantum security of (ring) signatures. In: Hanaoka, G., Shikata, J., Watanabe, Y. (eds.) PKC 2022, Part II. LNCS, vol. 13178, pp. 407–436. Springer, Cham (Mar 2022). https://doi.org/10.1007/978-3-030-97131-1_14
18. Chatterjee, R., Garg, S., Hajiabadi, M., Khurana, D., Liang, X., Malavolta, G., Pandey, O., Shiehian, S.: Compact ring signatures from learning with errors. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021, Part I. LNCS, vol. 12825, pp. 282–312. Springer, Cham, Virtual Event (Aug 2021). https://doi.org/10.1007/978-3-030-84242-0_11
19. Chiesa, A., Manohar, P., Spooner, N.: Succinct arguments in the quantum random oracle model. In: Hofheinz, D., Rosen, A. (eds.) TCC 2019, Part II. LNCS, vol. 11892, pp. 1–29. Springer, Cham (Dec 2019). https://doi.org/10.1007/978-3-030-36033-7_1
20. Chung, K.M., Fehr, S., Huang, Y.H., Liao, T.N.: On the compressed-oracle technique, and post-quantum security of proofs of sequential work. In: Canteaut, A., Standaert, F.X. (eds.) EUROCRYPT 2021, Part II. LNCS, vol. 12697, pp. 598–629. Springer, Cham (Oct 2021). https://doi.org/10.1007/978-3-030-77886-6_21
21. Chung, K.M., Hsieh, Y.C., Huang, M.Y., Huang, Y.H., Lange, T., Yang, B.Y.: Group signatures and accountable ring signatures from isogeny-based assumptions. Cryptology ePrint Archive, Report 2021/1368 (2021), https://eprint.iacr.org/2021/1368
22. Collins, D., Huguenin-Dumittan, L., Nguyen, N.K., Rolin, N., Vaudenay, S.: K-waay: Fast and deniable post-quantum X3DH without ring signatures. In: Balzarotti, D., Xu, W. (eds.) USENIX Security 2024. USENIX Association (Aug 2024), https://www.usenix.org/conference/usenixsecurity24/presentation/collins
23. Derler, D., Ramacher, S., Slamanig, D.: Post-quantum zero-knowledge proofs for accumulators with applications to ring signatures from symmetric-key primitives. In: Lange, T., Steinwandt, R. (eds.) Post-Quantum Cryptography - 9th International Conference, PQCrypto 2018. pp. 419–440. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-79063-3_20
24. Deutsch, D., Jozsa, R.: Rapid solution of problems by quantum computation. Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences **439**(1907), 553–558 (1992)
25. Devevey, J., Fallahpour, P., Passelègue, A., Stehlé, D.: A detailed analysis of Fiat-Shamir with aborts. In: Handschuh, H., Lysyanskaya, A. (eds.) CRYPTO 2023, Part V. LNCS, vol. 14085, pp. 327–357. Springer, Cham (Aug 2023). https://doi.org/10.1007/978-3-031-38554-4_11
26. Don, J., Fehr, S., Majenz, C.: The measure-and-reprogram technique 2.0: Multi-round Fiat-Shamir and more. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020, Part III. LNCS, vol. 12172, pp. 602–631. Springer, Cham (Aug 2020). https://doi.org/10.1007/978-3-030-56877-1_21

27. Don, J., Fehr, S., Majenz, C., Schaffner, C.: Security of the Fiat-Shamir transformation in the quantum random-oracle model. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part II. LNCS, vol. 11693, pp. 356–383. Springer, Cham (Aug 2019). https://doi.org/10.1007/978-3-030-26951-7_13
28. Don, J., Fehr, S., Majenz, C., Schaffner, C.: Efficient NIZKs and signatures from commit-and-open protocols in the QROM. In: Dodis, Y., Shrimpton, T. (eds.) CRYPTO 2022, Part II. LNCS, vol. 13508, pp. 729–757. Springer, Cham (Aug 2022). https://doi.org/10.1007/978-3-031-15979-4_25
29. Don, J., Fehr, S., Majenz, C., Schaffner, C.: Online-extractability in the quantum random-oracle model. In: Dunkelman, O., Dziembowski, S. (eds.) EUROCRYPT 2022, Part III. LNCS, vol. 13277, pp. 677–706. Springer, Cham (May / Jun 2022). https://doi.org/10.1007/978-3-031-07082-2_24
30. van Erven, T., Harremoes, P.: Rényi divergence and kullback-leibler divergence. IEEE Transactions on Information Theory **60**(7), 3797–3820 (Jul 2014). https://doi.org/10.1109/tit.2014.2320500, http://dx.doi.org/10.1109/TIT.2014.2320500
31. Fallahpour, P., Fehr, S., Huang, Y.H.: Tighter quantum security for Fiat-Shamir-with-aborts and hash-and-sign-with-retry signatures. Cryptology ePrint Archive, Report 2025/985 (2025), https://eprint.iacr.org/2025/985
32. Fiat, A., Shamir, A.: How to prove yourself: Practical solutions to identification and signature problems. In: Odlyzko, A.M. (ed.) CRYPTO'86. LNCS, vol. 263, pp. 186–194. Springer, Berlin, Heidelberg (Aug 1987). https://doi.org/10.1007/3-540-47721-7_12
33. Fouque, P.A., Hoffstein, J., Kirchner, P., Lyubashevsky, V., Pornin, T., Prest, T., Ricosset, T., Seiler, G., Whyte, W., Zhang, Z., et al.: Falcon: Fast-fourier lattice-based compact signatures over ntru. Submission to the NIST's post-quantum cryptography standardization process **36**(5), 1–75 (2018)
34. Gajland, P., Janneck, J., Kiltz, E.: A closer look at falcon. Cryptology ePrint Archive, Report 2024/1769 (2024), https://eprint.iacr.org/2024/1769
35. Gajland, P., Janneck, J., Kiltz, E.: Ring signatures for deniable AKEM: Gandalf's fellowship. In: Reyzin, L., Stebila, D. (eds.) CRYPTO 2024, Part I. LNCS, vol. 14920, pp. 305–338. Springer, Cham (Aug 2024). https://doi.org/10.1007/978-3-031-68376-3_10
36. Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Ladner, R.E., Dwork, C. (eds.) 40th ACM STOC. pp. 197–206. ACM Press (May 2008). https://doi.org/10.1145/1374376.1374407
37. Giacomelli, I., Madsen, J., Orlandi, C.: ZKBoo: Faster zero-knowledge for Boolean circuits. In: Holz, T., Savage, S. (eds.) USENIX Security 2016. pp. 1069–1083. USENIX Association (Aug 2016), https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/giacomelli
38. Grilo, A.B., Hövelmanns, K., Hülsing, A., Majenz, C.: Tight adaptive reprogramming in the QROM. In: Tibouchi, M., Wang, H. (eds.) ASIACRYPT 2021, Part I. LNCS, vol. 13090, pp. 637–667. Springer, Cham (Dec 2021). https://doi.org/10.1007/978-3-030-92062-3_22
39. Grover, L.K.: A fast quantum mechanical algorithm for database search. In: 28th ACM STOC. pp. 212–219. ACM Press (May 1996). https://doi.org/10.1145/237814.237866
40. Haque, A., Scafuro, A.: Threshold ring signatures: New definitions and post-quantum security. In: Kiayias, A., Kohlweiss, M., Wallden, P., Zikas, V. (eds.)

PKC 2020, Part II. LNCS, vol. 12111, pp. 423–452. Springer, Cham (May 2020). https://doi.org/10.1007/978-3-030-45388-6_15
41. Hashimoto, K., Katsumata, S., Kwiatkowski, K., Prest, T.: An efficient and generic construction for Signal's handshake (X3DH): Post-quantum, state leakage secure, and deniable. In: Garay, J. (ed.) PKC 2021, Part II. LNCS, vol. 12711, pp. 410–440. Springer, Cham (May 2021). https://doi.org/10.1007/978-3-030-75248-4_15
42. Hashimoto, K., Katsumata, S., Wiggers, T.: Bundled authenticated key exchange: A concrete treatment of (post-quantum) signal's handshake protocol. Cryptology ePrint Archive, Paper 2025/040 (2025), https://eprint.iacr.org/2025/040
43. Hoffstein, J., Pipher, J., Silverman, J.H.: NTRU: A ring-based public key cryptosystem. In: Third Algorithmic Number Theory Symposium (ANTS). LNCS, vol. 1423, pp. 267–288. Springer (Jun 1998)
44. Ishai, Y., Kushilevitz, E., Ostrovsky, R., Sahai, A.: Zero-knowledge from secure multiparty computation. In: Johnson, D.S., Feige, U. (eds.) 39th ACM STOC. pp. 21–30. ACM Press (Jun 2007). https://doi.org/10.1145/1250790.1250794
45. Katsumata, S., Niot, G., Tucker, I., Wiggers, T.: Comprehensive deniability analysis of signal handshake protocols: X3DH, PQXDH to fully post-quantum with deniable ring signatures. Cryptology ePrint Archive, Paper 2025/1090 (2025), https://eprint.iacr.org/2025/1090
46. Kret, E., Schmidt, R.: The pqxdh key agreement protocol. Online (2023), https://signal.org/docs/specifications/pqxdh/
47. Langlois, A., Stehlé, D., Steinfeld, R.: GGHLite: More efficient multilinear maps from ideal lattices. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 239–256. Springer, Berlin, Heidelberg (May 2014). https://doi.org/10.1007/978-3-642-55220-5_14
48. Li, J., Jiang, H., Wang, H., Duan, Q.: Tighter security proof of falcon in the quantum random oracle model. In: Chen, R., Deng, R.H., Yung, M. (eds.) Information Security and Cryptology. pp. 142–165. Springer Nature Singapore, Singapore (2026)
49. Lu, X., Au, M.H., Zhang, Z.: Raptor: A practical lattice-based (linkable) ring signature. In: Deng, R.H., Gauthier-Umaña, V., Ochoa, M., Yung, M. (eds.) ACNS 2019. LNCS, vol. 11464, pp. 110–130. Springer, Cham (Jun 2019). https://doi.org/10.1007/978-3-030-21568-2_6
50. Lyubashevsky, V.: Lattice signatures without trapdoors. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 738–755. Springer, Berlin, Heidelberg (Apr 2012). https://doi.org/10.1007/978-3-642-29011-4_43
51. Lyubashevsky, V., Nguyen, N.K., Seiler, G.: SMILE: Set membership from ideal lattices with applications to ring signatures and confidential transactions. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021, Part II. LNCS, vol. 12826, pp. 611–640. Springer, Cham, Virtual Event (Aug 2021). https://doi.org/10.1007/978-3-030-84245-1_21
52. Marlinspike, M., Perrin, T.: The x3dh key agreement protocol. Online (2016), https://signal.org/docs/specifications/x3dh/
53. Micciancio, D., Regev, O.: Worst-case to average-case reductions based on Gaussian measures. In: 45th FOCS. pp. 372–381. IEEE Computer Society Press (Oct 2004). https://doi.org/10.1109/FOCS.2004.72
54. Mittal, B.: Ring trapdoor functions: A lattice-based framework for secure ring signatures, https://bhumikamittal.in/assets/pdf/undergrad_thesis_bhumika.pdf

55. Nielsen, M.A., Chuang, I.L.: Quantum computation and quantum information. Cambridge university press (2010)

56. Niot, G.: Practical deniable post-quantum X3DH: A lightweight split-KEM for k-waay. Cryptology ePrint Archive, Paper 2025/853 (2025), https://eprint.iacr.org/2025/853

57. Peikert, C., Rosen, A.: Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In: Halevi, S., Rabin, T. (eds.) TCC 2006. LNCS, vol. 3876, pp. 145–166. Springer, Berlin, Heidelberg (Mar 2006). https://doi.org/10.1007/11681878_8

58. Pöppelmann, T., Ducas, L., Güneysu, T.: Enhanced lattice-based signatures on reconfigurable hardware. In: Batina, L., Robshaw, M. (eds.) CHES 2014. LNCS, vol. 8731, pp. 353–370. Springer, Berlin, Heidelberg (Sep 2014). https://doi.org/10.1007/978-3-662-44709-3_20

59. Prest, T.: Gaussian sampling in lattice-based cryptography. Theses, Ecole normale supérieure - ENS PARIS (Dec 2015), https://theses.hal.science/tel-01245066

60. Rivest, R.L., Shamir, A., Tauman, Y.: How to leak a secret. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 552–565. Springer, Berlin, Heidelberg (Dec 2001). https://doi.org/10.1007/3-540-45682-1_32

61. Stehlé, D., Steinfeld, R.: Making NTRU as secure as worst-case problems over ideal lattices. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 27–47. Springer, Berlin, Heidelberg (May 2011). https://doi.org/10.1007/978-3-642-20465-4_4

62. Wang, J., Sun, B.: Ring signature schemes from lattice basis delegation. In: Qing, S., Susilo, W., Wang, G., Liu, D. (eds.) ICICS 11. LNCS, vol. 7043, pp. 15–28. Springer, Berlin, Heidelberg (Nov 2011). https://doi.org/10.1007/978-3-642-25243-3_2

63. Yuen, T.H., Esgin, M.F., Liu, J.K., Au, M.H., Ding, Z.: DualRing: Generic construction of ring signatures with efficient instantiations. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021, Part I. LNCS, vol. 12825, pp. 251–281. Springer, Cham, Virtual Event (Aug 2021). https://doi.org/10.1007/978-3-030-84242-0_10

64. Zhandry, M.: How to construct quantum random functions. In: 53rd FOCS. pp. 679–687. IEEE Computer Society Press (Oct 2012). https://doi.org/10.1109/FOCS.2012.37

65. Zhandry, M.: Secure identity-based encryption in the quantum random oracle model. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 758–775. Springer, Berlin, Heidelberg (Aug 2012). https://doi.org/10.1007/978-3-642-32009-5_44

66. Zhandry, M.: How to record quantum queries, and applications to quantum indifferentiability. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part II. LNCS, vol. 11693, pp. 239–268. Springer, Cham (Aug 2019). https://doi.org/10.1007/978-3-030-26951-7_9

67. Zhang, X., Li, Z., Steinfeld, R., Zhao, R.K., Liu, J.K., Yuen, T.H.: Pegasus and pegaring: Efficient (ring) signatures from sigma-protocols for power residue prfs with (q) rom security. Cryptology ePrint Archive (2025)

A Additional Preliminaries

A.1 Probabilities

Let P and Q be probability distributions on the same countable alphabet Ω. Then, the statistical distance between P and Q is $\Delta(P,Q) := \sum_{x\in\Omega} |P(x) - Q(x)|/2$.

The Bhattacharyya coefficient of two probability distributions P, Q on a countable alphabet Ω is

$$\mathrm{BC}(P,Q) := \sum_{x\in\Omega} \sqrt{P(x)Q(x)}.$$

The Hellinger distance is defined via

$$\mathrm{dh}(P,Q)^2 := 1 - \mathrm{BC}(P,Q) = \frac{1}{2}\sum_{x\in\Omega}\left(\sqrt{P(x)} - \sqrt{Q(x)}\right)^2.$$

It is bounded by the statistical distance as $\mathrm{dh}(P,Q) \leq \sqrt{\Delta(P,Q)}$.

The Rényi divergence of order $\alpha \in (1,\infty)$ of distributions P, Q with $\mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q)$ is

$$R_\alpha(P\|Q) = \left(\sum\nolimits_{x\in\mathrm{Supp}(P)} P(x)^\alpha \cdot Q(x)^{1-\alpha}\right)^{\frac{1}{\alpha-1}}$$

and the Rényi divergence of order ∞ is defined as

$$R_\infty(P\|Q) = \sup\nolimits_{x\in\mathrm{Supp}(P)} P(x)/Q(x) \in [1,\infty].$$

These are the multiplicative relations often used in cryptography. The divergences are defined as

$$D_\alpha(P\|Q) := \ln(R_\alpha(P\|Q)).$$

The Kullback-Leibler divergence for two distributions over the same countable set Ω is defined as

$$D_{KL}(P\|Q) = \sum_{x\in\Omega} \ln\left(\frac{P(x)}{Q(x)}\right)\cdot P(x)$$

with the convention that $\ln(x/0) = +\infty$ for any $x > 0$ and 0 for $x = 0$. Furthermore, Pinsker's inequality gives an implicit bound on the statistical distance $\Delta(P,Q) \leq \sqrt{D_{KL}(P\|Q)/2}$. As $D_{KL}(P\|Q) \leq D_\alpha(P\|Q)$ for all $\alpha > 1$, the Rényi divergence can be used to bound the statistical distance.

A.2 Signatures

We recall the syntax of digital signatures.

Definition 4 (Signatures). *A signature scheme* Sig *is defined as a tuple* $(\mathsf{Gen}, \mathsf{Sign}, \mathsf{Vf})$ *of the following three algorithms:*

- $\mathsf{Gen}(1^\lambda) \to (\mathsf{sk}, \mathsf{pk})$*: The probabilistic key generation algorithm returns a pair of secret key* sk *and public key* pk*, where the public key defines the message space.*
- $\mathsf{Sign}(\mathsf{sk}, m) \to \sigma$*: The probabilistic signature algorithm takes in a secret key and a message* m*, and outputs a signature* σ*.*
- $\mathsf{Vf}(\mathsf{pk}, m, \sigma) \to b$*: The deterministic verification algorithm takes in the public key* pk*, a message* m *and a signature* σ*. It outputs either* false *(*$b = 0$*) or* true *(*$b = 1$*).*

Sig *has correctness error* ϵ *if for all* $(\mathsf{pk}, \mathsf{sk}) \in \mathrm{Supp}(\mathsf{Gen})$ *and any* m*:*

$$\Pr[\mathsf{Vf}(\mathsf{pk}, m, \mathsf{Sign}(\mathsf{sk}, m)) \neq 1] \leq \epsilon,$$

where the probability is taken over the randomness of Sign*.*

Definition 5 (SUF-CMA of Signatures). *Consider the unforgeability game in Fig. 3. We say that a signature* Sig *is strongly unforgeable under* q_s *chosen message attacks, if for any ppt adversary* $\mathcal{A}$

$$\mathsf{Adv}^{\text{SUF-CMA}}_{\mathcal{A},q_s}(\lambda) := \Pr[1 \leftarrow \text{SUF-CMA}_{\mathcal{A},\mathsf{Sig},q_s}] = \mathsf{negl}(\lambda).$$

$\text{SUF-CMA}_{\mathcal{A},\mathsf{Sig},q_s}$	$\mathsf{OSign}(m)$
$q \leftarrow 0; \mathcal{L} \leftarrow \emptyset$	**if** $q = q_s$:
$(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}(1^\lambda)$	**return** $\perp$
$(m^*, \sigma^*) \leftarrow \mathcal{A}^{\mathsf{OSign}}(\mathsf{pk})$	$q \leftarrow q + 1$
if $(m^*, \sigma^*) \in \mathcal{L}$:	$\sigma \leftarrow \mathsf{Sign}(\mathsf{sk}, m)$
return 0 // return false	$\mathcal{L} \leftarrow \mathcal{L} \cup \{(m, \sigma)\}$
return $\mathsf{Vf}(\mathsf{pk}, m^*, \sigma^*)$	**return** σ

Fig. 3. Unforgeability games for signatures.

A.3 Ring Signatures

Definition 6 (Ring Signatures). *A ring signature scheme (RSS)* RSig *is a quadruple of ppt algorithms* $(\mathsf{Stp}, \mathsf{KGen}, \mathsf{Sign}, \mathsf{Vf})$ *such that:*

- $\mathsf{Stp}(1^\lambda) \to \mathsf{p}$*: The setup algorithm takes as input the security parameter* 1^λ *and outputs public parameters* p*. These parameters also define the message space and an upper bound on the ring size* κ*.*

- $\mathsf{KGen}(\mathsf{p}) \to (\mathsf{pk}, \mathsf{sk})$*: The key generation algorithm takes as input the public parameters* p *and produces a pair of public and private keys* $(\mathsf{pk}, \mathsf{sk})$.
- $\mathsf{Sign}(\mathsf{sk}, \rho, m) \to \sigma$*: The signing algorithm takes as input the secret key of the signer* sk*, a list of public keys* $\rho = \{\mathsf{pk}_1, \ldots, \mathsf{pk}_N\}$ *defining the ring and a message* m*. The ring must satisfy the size bound, i.e.,* $N \leq \kappa$*, and* $\exists i \in [N] : \mu(\mathsf{sk}) = \mathsf{pk}_i$.[9] *It outputs a signature* σ.
- $\mathsf{Vf}(\rho, m, \sigma) \to b$*: The deterministic verification algorithm takes as input a list of public keys* $\rho = \{\mathsf{pk}_1, \ldots, \mathsf{pk}_N\}$ *with* $N \leq \kappa$*, a message* m *and a signature* σ*. It outputs either* false *(*$b = 0$*) or* true *(*$b = 1$*).*

RSig *is* $\delta(\kappa)$*-correct if for any* $\mathsf{p} \leftarrow \mathsf{Stp}(1^\lambda)$*, all* $N \leq \kappa$*, any* $\{(\mathsf{pk}_j, \mathsf{sk}_j)\}_{j \in [N]} \subseteq \mathrm{Supp}(\mathsf{KGen}(\mathsf{p}))$ *defining* $\rho = \{\mathsf{pk}_1, \ldots, \mathsf{pk}_N\}$*, any* m *and any* $i \in [N]$

$$\Pr[\mathsf{Vf}(\rho, m, \mathsf{Sign}(\mathsf{sk}_i, \rho, m)) \neq 1] \leq \delta(\kappa).$$

Unforgeability of Ring Signatures. As in standard EUF-CMA, a forgery adversary against a RSS must output a forgery σ^* for a message m^* for a ring $\rho^* \subseteq \{\mathsf{pk}_1, \ldots, \mathsf{pk}_N\}$ of valid public keys. The adversary can make adaptive queries to a signing oracle with a message, a signer position, and a list of public keys. The signer position must refer to a valid public key. This is referred to as "insider security" [9]. In the following, we treat the size bound for signing queries and verification queries as implicit, i.e. the ring size of queries cannot exceed κ.

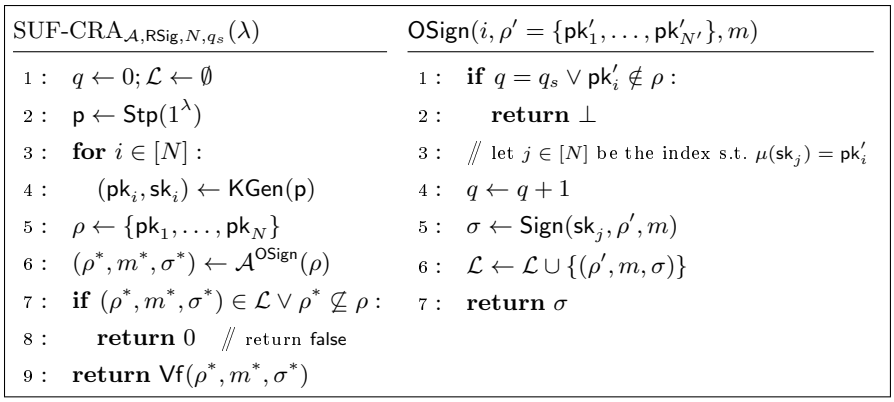

Fig. 4. Unforgeability game for ring signatures.

Definition 7 (UF-NRA of Ring Signatures). *A RSS* RSig *is unforgeable under no-ring attacks for* $N \leq \kappa$ *if, for any ppt adversary* $\mathcal{A}$

$$\mathsf{Adv}^{\text{UF-NRA}}_{\mathcal{A},\mathsf{RSig},N}(\lambda) := \Pr[1 \leftarrow \text{UF-NRA}_{\mathcal{A},\mathsf{RSig},N}(\lambda)] = \mathsf{negl}(\lambda),$$

[9] We assume (w.l.o.g.) that there exists a one-way function μ such that $\mu(\mathsf{sk}) = \mathsf{pk}$ for all $(\mathsf{pk}, \mathsf{sk}) \in \mathrm{Supp}(\mathsf{KGen})$.

where $\text{UF-NRA}_{\mathcal{A},\mathsf{RSig},N} = \text{SUF-CRA}_{\mathcal{A},\mathsf{RSig},N,0}$.

Definition 8 (SUF-CRA of Ring Signatures). *Consider the unforgeability game in Fig. 4. We say a RSS* RSig *is strongly unforgeable under q_s chosen ring attacks for $N \leq \kappa$ if, for any ppt adversary $\mathcal{A}$*

$$\mathsf{Adv}^{\text{SUF-CRA}}_{\mathcal{A},\mathsf{RSig},N,q_s}(\lambda) := \Pr[1 \leftarrow \text{SUF-CRA}_{\mathcal{A},\mathsf{RSig},N,q_s}(\lambda)] = \mathsf{negl}(\lambda).$$

Consider the unforgeability game in Fig. 4. Define the UF-CRA1 game as the SUF-CRA game, except that i) $\mathcal{L}$ contains pairs (ρ, m) (without the signatures), ii) a forgery counts as fresh if it is for a pair $(\rho^*, m^*) \notin \mathcal{L}$ and iii) the signing oracle cannot be queried for (ρ', m) already in $\mathcal{L}$.

Definition 9 (UF-CRA1 of Ring Signatures). *We say a RSS* RSig *is one-per-message weakly unforgeable under q_s chosen ring attacks for $N \leq \kappa$ if, for any ppt adversary $\mathcal{A}$*

$$\mathsf{Adv}^{\text{UF-CRA1}}_{\mathcal{A},\mathsf{RSig},N,q_s}(\lambda) := \Pr[1 \leftarrow \text{UF-CRA1}_{\mathcal{A},\mathsf{RSig},N,q_s}(\lambda)] = \mathsf{negl}(\lambda).$$

Anonymity of Ring Signatures. We consider *anonymity under full key exposure* introduced in [9]. As in [15,35], we parameterize the anonymity notion with q_c, defining the number of allowed calls to the challenge oracle.

$\text{ANON}_{\mathcal{A},\mathsf{RSig},N,q_c}(\lambda)$	$\mathsf{CHAL}(i_0, i_1, \rho' = \{\mathsf{pk}'_1, \ldots, \mathsf{pk}'_{N'}\}, m)$
$q \leftarrow 0; b \leftarrow \{0,1\}$	**if** $q = q_c \vee \mathsf{pk}'_{i_0} \notin \rho \vee \mathsf{pk}'_{i_1} \notin \rho$:
$\mathsf{p} \leftarrow \mathsf{Stp}(1^\lambda)$	**return** $\perp$
for $i \in [N]$:	// let $j \in [N]$ be the index s.t. $\mu(\mathsf{sk}_j) = \mathsf{pk}'_{i_b}$
$(\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow \mathsf{KGen}(\mathsf{p})$	$q \leftarrow q + 1$
$\rho \leftarrow \{\mathsf{pk}_i\}_{i\in[N]}; \mathsf{SK} \leftarrow \{\mathsf{sk}_i\}_{i\in[N]}$	**return** $\mathsf{Sign}(\mathsf{sk}_j, \rho', m)$
$b' \leftarrow \mathcal{A}^{\mathsf{CHAL}}(\rho, \mathsf{SK})$	
return $b = b'$	

Fig. 5. The anonymity game for RSSs under full key exposure.

Definition 10 (Anonymity of Ring Signatures). *Consider the anonymity game in Fig. 5. We say a RSS* RSig *is anonymous under full key exposure with q_c challenge queries if, for $N \leq \kappa$ and any ppt adversary $\mathcal{A}$*

$$\mathsf{Adv}^{\text{ANON}}_{\mathcal{A},\mathsf{RSig},N,q_c}(\lambda) := |\Pr[\text{ANON}_{\mathcal{A},\mathsf{RSig},N,q_c}(\lambda) \Rightarrow \mathsf{true}] - 1/2| = \mathsf{negl}(\lambda)$$

A.4 Sigma Protocols

Definition 11 (Σ-protocol). *A Σ-protocol $\Sigma = (\mathsf{Gen}, \mathsf{P}, \mathsf{V})$ is a three-move identification protocol consisting of five ppt algorithms $\mathsf{Gen}, \mathsf{P}_1, \mathsf{P}_2, \mathsf{V}_1, \mathsf{V}_2$ where V_2 is deterministic. We assume $\mathsf{P}_1, \mathsf{P}_2$ to share state, and likewise V_1 and V_2. Denote the challenge space by $\mathcal{C}$. Then, Σ behaves as follows.*

1. *The prover gets a pair of instance and witness $(\mathsf{inst}, w) \leftarrow \mathsf{Gen}(1^\lambda)$ where $w \in \mathcal{W}$ and publishes its instance $\mathsf{inst} \in \mathcal{I}$.*[10]
2. *The prover runs P_1 with input inst to get a commitment $(\mathsf{com}, \mathsf{state}) \leftarrow \mathsf{P}_1(\mathsf{inst})$, and sends com to the verifier.*
3. *The verifier runs V_1 with input $\mathsf{inst}, \mathsf{com}$, and generates an independent random challenge from the challenge set $\mathcal{C}$ and sends it to the prover.*
4. *The prover runs P_2 with input $\mathsf{state}, w, \mathsf{ch}$ to compute a response $\mathsf{rsp} \leftarrow \mathsf{P}_2(\mathsf{state}, w, \mathsf{ch})$ for the commitment com, and sends it to the verifier.*
5. *The verifier runs V_2 with input $\mathsf{inst}, \mathsf{com}, \mathsf{ch}, \mathsf{rsp}$ and outputs either false or true.*

A protocol transcript $(\mathsf{com}, \mathsf{ch}, \mathsf{rsp})$ is said to be valid in case $\mathsf{V}_2(\mathsf{inst}, \mathsf{com}, \mathsf{ch}, \mathsf{rsp})$ outputs true. We say a Σ-protocol is complete, if

$$\Pr\left[\mathsf{V}_2(\mathsf{inst}, \mathsf{com}, \mathsf{ch}, \mathsf{rsp}) : \begin{smallmatrix}(\mathsf{inst},w)\leftarrow\mathsf{Gen}(1^\lambda);\ (\mathsf{com},\mathsf{state})\leftarrow\mathsf{P}_1(\mathsf{inst});\\ \mathsf{ch}\leftarrow\mathsf{V}_1(\mathsf{inst},\mathsf{com});\ \mathsf{rsp}\leftarrow\mathsf{P}_2(\mathsf{state},w,\mathsf{ch})\end{smallmatrix}\right] = 1.$$

If a response and a challenge uniquely define the commitment, then we call the Σ-protocol *commitment-recoverable*.

Security for the Verifier. The verifier needs to be convinced that the prover knows the witness. Therefore, an adversary without knowing the witness should not be able to impersonate an honest prover.

Definition 12 (Impersonation under Passive Attacks). *We say that a Σ-protocol Σ is secure against impersonations under passive attacks, if for any ppt $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ and honest verifier $\mathsf{V} = (\mathsf{V}_1, \mathsf{V}_2)$, we have*

$$\mathsf{Adv}^{\mathrm{IMP}}_{\mathcal{A},\Sigma}(\lambda) := \Pr\left[\mathsf{V}_2(\mathsf{inst}, \mathsf{com}, \mathsf{ch}, \mathsf{rsp}) : \begin{smallmatrix}(\mathsf{inst},w)\leftarrow\Sigma.\mathsf{Gen}(1^\lambda);\\ (\mathsf{com},\mathsf{state})\leftarrow\mathcal{A}_1(\mathsf{inst});\\ \mathsf{ch}\leftarrow\mathsf{V}_1(\mathsf{inst},\mathsf{com});\\ \mathsf{rsp}\leftarrow\mathcal{A}_2(\mathsf{state},\mathsf{ch})\end{smallmatrix}\right] = \mathsf{negl}(\lambda).$$

An important security notion for Σ-protocols is *special soundness*.

Definition 13 (Special Soundness). *We say that a Σ-protocol Σ is* special sound *if there exists an ppt* Ext *that on input of any* inst *from* $(\mathsf{inst}, w) \leftarrow \mathsf{Gen}(1^\lambda)$ *and any two valid transcripts* $(\mathsf{com}, \mathsf{ch}_0, \mathsf{rsp}_0)$ *and* $(\mathsf{com}, \mathsf{ch}_1, \mathsf{rsp}_1)$ *with* $\mathsf{ch}_0 \neq \mathsf{ch}_1$, *outputs some witness w' such that $(\mathsf{inst}, w') \in R$.*

[10] We assume (w.l.o.g.) that there exists a function μ such that $\mu(w) = \mathsf{inst}$ for all $(\mathsf{inst}, w) \in \mathrm{Supp}(\mathsf{Gen})$. If this is not the case, we can re-define the witness to contain a copy of the instance.

A Σ-protocol has unique responses if there is only one valid response for a commitment-challenge-pair. A relaxation is computational unique responses.

Definition 14 (Computational Unique Responses). *We say that a Σ-protocol Σ has* computational unique responses *if for any ppt $\mathcal{A}$, we have*

$$\mathsf{Adv}^{\mathrm{CUR}}_{\mathcal{A},\Sigma}(\lambda) := \Pr\left[\begin{array}{c}\mathsf{rsp}_0 \neq \mathsf{rsp}_1 \\ \wedge \mathsf{V}_2(\mathsf{inst},\mathsf{com},\mathsf{ch},\mathsf{rsp}_0) \\ \wedge \mathsf{V}_2(\mathsf{inst},\mathsf{com},\mathsf{ch},\mathsf{rsp}_1)\end{array} : \begin{array}{c}(\mathsf{inst},w)\leftarrow\mathsf{Gen}(1^\lambda); \\ (\mathsf{com},\mathsf{ch},\mathsf{rsp}_0,\mathsf{rsp}_1)\leftarrow\mathcal{A}(\mathsf{inst})\end{array}\right] = \mathsf{negl}(\lambda).$$

Security for the Prover. The prover wants to ensure that the witness is not revealed to others. There are two levels to this. First, the instance should not reveal anything about the witness, and secondly, the commitments and responses should not reveal anything about the witness.

Definition 15 (Witness Recovery). *We say that a Σ-protocol Σ for a witness relation R is secure against witness recovery, if for any ppt $\mathcal{A}$, we have*

$$\mathsf{Adv}^{\text{W-REC}}_{\mathcal{A},\Sigma}(\lambda) := \Pr\left[(\mathsf{inst},w') \in R : \begin{array}{c}(\mathsf{inst},w)\leftarrow\mathsf{Gen}(1^\lambda); \\ w'\leftarrow\mathcal{A}(\mathsf{inst})\end{array}\right] = \mathsf{negl}(\lambda).$$

Definition 16 (Honest Verifier Zero Knowledge). *We say that a Σ-protocol Σ is* statistically *special $\epsilon^{\mathrm{HVZK}}_\Sigma$ HVZK if there exists a ppt simulator, such that for any $\lambda, (\mathsf{inst},w) \leftarrow \mathsf{Gen}(1^\lambda)$ and any challenge $\mathsf{ch} \in \mathcal{C}$, we have*

$$\Delta(P_{\mathsf{ch}}(\mathsf{inst},w), \mathsf{Sim}(\mathsf{inst},\mathsf{ch})) \leq \epsilon^{\mathrm{HVZK}}_\Sigma$$

where $P_{\mathsf{ch}}(\mathsf{inst},w)$ denotes the distribution of honestly generated transcripts according to the algorithm's description, conditioned on having challenge ch. Furthermore, we call it $\delta^{\alpha\text{-HVZK}}_\Sigma$-divergence HVZK for $\alpha \in (1,\infty]$ if

$$R_\alpha(P_{\mathsf{ch}}(\mathsf{inst},w) \| \mathsf{Sim}(\mathsf{inst},\mathsf{ch})) \leq \delta^{\alpha\text{-HVZK}}_\Sigma.$$

We also define the Kullback-Leibler divergence for HVZK as

$$D_{KL}(P_{\mathsf{ch}}(\mathsf{inst},w) \| \mathsf{Sim}(\mathsf{inst},\mathsf{ch})) \leq \epsilon^{KL\text{-HVZK}}_\Sigma.$$

A simulator produces commitments and responses. However, the distribution of commitments in the range of all commitments is not immediate, so we consider the notion of commitment min-entropy for a HVZK simulator for *any* input.

Definition 17 (Simulator Commitment Min-Entropy). *We say that simulator Sim for Σ has $\beta(\lambda)$ commitment min-entropy, if*

$$\max_{C,\mathsf{ch}\in\mathcal{C},\mathsf{inst}\in\{0,1\}^*} \Pr[C = \mathsf{com} : (\mathsf{com},\mathsf{rsp}) \leftarrow \mathsf{Sim}(\mathsf{inst},\mathsf{ch})] \leq 2^{-\beta(\lambda)}.$$

For protocols with statistical HVZK, the notion is equivalent to the more common notion of commitment min-entropy.

A.5 Quantum Computation Basics

The length of a vector $|\psi\rangle \in \mathbb{C}^d$ can be measured with the Euclidean norm

$$\||\psi\rangle\|_2 := \sqrt{\langle\psi|\psi\rangle}.$$

Given an operator $M \in \mathcal{L}(\mathbb{C}^d)$, define its trace norm and operator norm as

$$\|M\|_{tr} := \operatorname{Tr}\sqrt{M^\dagger M} \quad \text{and} \quad \|M\|_\infty := \max_{\text{unit vectors } |\psi\rangle \in \mathbb{C}^d} \|M|\psi\rangle\|_2.$$

For a vector $|\psi\rangle \in \mathbb{C}^d$ and two unitaries U, V on $\mathbb{C}^d$ it holds that

$$\left\| U^\dagger |\psi\rangle\langle\psi| U - V^\dagger |\psi\rangle\langle\psi| V \right\|_{tr} \leq 2\|U - V\|_\infty. \tag{7}$$

Both the trace and the operator norm are also invariant under multiplication with unitaries. Any rank-1 operator $M = |\psi\rangle\langle\phi| \in \mathcal{L}(\mathbb{C}^d)$ satisfies

$$\|M\|_\infty = \||\psi\rangle\|_2 \||\phi\rangle\|_2. \tag{8}$$

Grover Search Let $f : \{0,1\}^n \to \{0,1\}$ be an n-bit function. An element $x \in \{0,1\}^n$ is marked, if and only if $f(x) = 1$. Grover's algorithm [39] allows for finding a marked element in time $\mathcal{O}\left(\sqrt{N/M}\right)$ where M is the number of marked elements and $N = 2^n$. We follow the notation of [55] in this work. The initial state of Grover's algorithm is an equal superposition $|\psi\rangle = \frac{1}{\sqrt{N}} \sum_{x\in\{0,1\}^n} |x\rangle$. By separating the marked and unmarked elements, the state can be written as

$$|\psi\rangle = \sqrt{\frac{N-M}{N}}\,|\alpha\rangle + \sqrt{\frac{M}{N}}\,|\beta\rangle\,,$$

where

$$|\alpha\rangle = \frac{1}{\sqrt{N-M}} \sum_{\substack{x\in\{0,1\}^n \\ f(x)=0}} |x\rangle\,, \qquad |\beta\rangle = \frac{1}{\sqrt{M}} \sum_{\substack{x\in\{0,1\}^n \\ f(x)=1}} |x\rangle\,.$$

This is interpreted geometrically by defining an angle θ such that $\cos(\theta/2) = \sqrt{(N-M)/N}$, which yields $|\psi\rangle = \cos\left(\frac{\theta}{2}\right)|\alpha\rangle + \sin\left(\frac{\theta}{2}\right)|\beta\rangle$. The Grover iteration G applied q times takes the state to

$$\left|\psi^{(q)}\right\rangle := G^q |\psi\rangle = \cos\left(\frac{2q+1}{2}\theta\right)|\alpha\rangle + \sin\left(\frac{2q+1}{2}\theta\right)|\beta\rangle\,. \tag{9}$$

A.6 NTRU Lattices

We introduce the lattice notation used in this appendix. Most definitions are adapted from [35,34], with details omitted where they are not essential.

We work with polynomial rings of the form $\mathcal{R} = \mathbb{Z}[X]/f(X)$ where $f(X) = X^M + 1$ and M is a power of two, i.e., $\exists k : M = 2^k$. Polynomials are denoted by lower-case bold letters. We write $\mathcal{R}_q = \mathbb{Z}_q[X]/f(X)$.

A lattice $\boldsymbol{\Lambda} \subseteq \mathbb{R}^n$ is defined by a basis $\boldsymbol{B} = \{\boldsymbol{b}_1, \ldots, \boldsymbol{b}_m\}$ of linearly independent vectors $\boldsymbol{b}_i \in \mathbb{R}^n$. They form the lattice

$$\boldsymbol{\Lambda}(\boldsymbol{B}) = \boldsymbol{B}\mathbb{Z}^m = \left\{ \sum_{i \in [m]} c_i \boldsymbol{b}_i : c_1, \ldots, c_m \in \mathbb{Z} \right\}.$$

For a basis $\boldsymbol{B}$, denote the Gram-Schmidt norm from [36] as $\|\boldsymbol{B}\|_{GS} = \max_{i \in [m]} \|\tilde{\boldsymbol{b}}_i\|$, where $\tilde{\boldsymbol{B}}$ is the Gram-Schmidt orthogonalization of $\boldsymbol{B}$.

A discrete Gaussian over a lattice $\boldsymbol{\Lambda}$ is defined via the n-dimensional Gaussian function $\rho_{s,\boldsymbol{c}} : \mathbb{R}^n \to (0, 1]$ on $\mathbb{R}^n$ centered at $\boldsymbol{c} \in \mathbb{R}^n$ with standard deviation $s > 0$:

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

A discrete Gaussian over $\boldsymbol{\Lambda}$ is the Gaussian function normalized by all points of the lattice

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

We may omit $\boldsymbol{c}$ if it is $\boldsymbol{0}$, and we may omit $\boldsymbol{\Lambda}$, if it is $\mathbb{Z}^M$. We write $\boldsymbol{f} \leftarrow \mathcal{D}_{\mathcal{R},s}$ to denote the polynomial $\boldsymbol{f}$ whose coefficient embedding[11] is sampled as $\mathcal{D}_{\mathbb{Z}^M,s}$.

W denote by $\eta_\epsilon(\boldsymbol{\Lambda})$ the smoothing parameter of a lattice $\boldsymbol{\Lambda}$ (introduced and properly defined in [53]). Intuitively, the smoothing parameter is the smallest s such that a discrete Gaussian distribution with standard deviation s looks "smooth" (nearly uniform), when it is wrapped over the quotient $\mathbb{R}^n/\boldsymbol{\Lambda}$.

Lemma 11 ([57, Lemma 2.11]). *For any n-dimensional lattice $\boldsymbol{\Lambda}$, center $\boldsymbol{c} \in \mathbb{R}^n$, $\epsilon > 0$, and $s \geq 2\eta_\epsilon(\boldsymbol{\Lambda})$, and for every $\boldsymbol{x} \in \Lambda$, we have*

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

In particular, for $\epsilon < 1/3$, the min-entropy is at least $n - 1$.

Lemma 12 ([50,4]). *Let $n > 1$ and $s > 0$.*

1. *For any $\tau > 0$, $\Pr[|z| > \tau s : z \leftarrow \mathcal{D}_{\mathbb{Z},s}] \leq 2e^{\frac{-\tau^2}{2}}$.*
2. *For any $\tau > 1$, $\Pr[\|\boldsymbol{z}\|_2 > \tau s\sqrt{n} : \boldsymbol{z} \leftarrow \mathcal{D}_{\mathbb{Z}^n,s}] \leq \tau^n e^{\frac{n}{2}(1-\tau^2)}$.*

Definition 18 (NTRU Lattice). *Let M be a power of two, q prime, $\boldsymbol{f}, \boldsymbol{g} \in \mathcal{R} = \mathbb{Z}[X]/(X^M + 1)$ and $\boldsymbol{h} = \boldsymbol{g}\boldsymbol{f}^{-1} \mod q$. The NTRU lattice of $\boldsymbol{h}$ is defined by the coefficient embedding of*

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

We denote this lattice as $\boldsymbol{\Lambda}_{\boldsymbol{h},q}$.

[11] This embedding identifies polynomials in $\mathcal{R}$ with their coefficient vectors in $\mathbb{Z}^M$.

The basis of the lattice can be computed easily from the anticirculant matrix of $\boldsymbol{h}$. The sampling procedure also generates $\boldsymbol{h}$ typically such that it is in $\mathcal{R}_q^*$, i.e., that it is coprime to q.

Lemma 13 (NTRU Trapdoor Generation [43,59]). *Let $\mathcal{R} = \mathbb{Z}[X]/(X^M+1)$. There exists a ppt algorithm, $\mathsf{TpdGen}_{\mathsf{NTRU}}(q,\alpha)$, that given a modulus q, and a target quality α, returns a public key $\boldsymbol{h} \in \mathcal{R}_q$, and the trapdoor $(\boldsymbol{f},\boldsymbol{g}) \in \mathcal{R}\times\mathcal{R}$ such that they each define a basis $\boldsymbol{B}_{\boldsymbol{h}}$ and $\boldsymbol{B}_{\boldsymbol{f},\boldsymbol{g}}$ for the same lattice. Furthermore, $\|\boldsymbol{B}_{\boldsymbol{f},\boldsymbol{g}}\|_{GS} \leq \alpha\sqrt{q}$.*

Corollary 2 ([34, Corollary 2]). *Let q be prime, $\boldsymbol{h} \in \mathcal{R}_q\backslash\{\boldsymbol{0}\}, \alpha \in (1,\infty), \epsilon \in (0,1/2), s \geq \eta_{\epsilon(\boldsymbol{\Lambda}_{\boldsymbol{h},q})}$ and Q be the distribution of $\boldsymbol{h}\boldsymbol{u}+\boldsymbol{v}$ where $\boldsymbol{u},\boldsymbol{v} \leftarrow \mathcal{D}_{\mathcal{R},s,\boldsymbol{0}}$. Then it holds that*

$$R_\alpha(\mathcal{U}(\mathcal{R}_q)\|Q) \lesssim 1 + \frac{2\alpha\epsilon^2}{(1-\epsilon)^2}.$$

Corollary 3 ([35, Corollary 1 that uses [59]]). *Let M be a positive integer, $\alpha > 1$ and $\epsilon \in (0,1/4)$. Then, there exists a preimage sampling algorithm $\mathsf{PreSmp}_{\mathsf{NTRU}}(\boldsymbol{B},s,\boldsymbol{c})$, such that for any basis $\boldsymbol{B}$, standard deviation $s \geq \eta_{\epsilon(\mathbb{Z}^{2M})}\cdot\|\boldsymbol{B}\|_{GS}$ and arbitrary syndrome $\boldsymbol{c}$, the Kullback-Leibler divergence and Rényi divergence are bounded as*

$$D_{KL}(\mathsf{PreSmp}_{\mathsf{NTRU}}(\boldsymbol{B},s,\boldsymbol{c})\|\mathcal{D}_{\boldsymbol{\Lambda}(\boldsymbol{B}),s,\boldsymbol{c}}) \lesssim 2\epsilon^2$$

and

$$R_\alpha(\mathsf{PreSmp}_{\mathsf{NTRU}}(\boldsymbol{B},s,\boldsymbol{c})\|\mathcal{D}_{\boldsymbol{\Lambda}(\boldsymbol{B}),s,\boldsymbol{c}}) \lesssim 1 + 2\alpha\epsilon^2.$$

We define two hardness problems for NTRU lattices with respect to the trapdoor algorithm $\mathsf{TpdGen}_{\mathsf{NTRU}}$.

Definition 19 (NTRU-SIS). *Let $\mathcal{R} := \mathbb{Z}[X]/(X^M+1)$. The* short integer solution *problem relative to the NTRU trapdoor algorithm $\mathsf{TpdGen}_{\mathsf{NTRU}}$ with parameters $m,q,\alpha,\beta > 0$ is defined via the $\mathcal{R}$-NTRU-SIS game from Fig. 6. We define the advantage of $\mathcal{A}$ in $\mathcal{R}$-NTRU-SIS as*

$$\mathsf{Adv}^{\mathcal{R}\text{-NTRU-SIS}}_{\mathcal{A},m,q,\alpha,\beta}(\lambda) := \Pr[1 \leftarrow \mathcal{R}\text{-NTRU-SIS}_{\mathcal{A},m,q,\alpha,\beta}].$$

Definition 20 (t-NTRU-ISIS). *Let $\mathcal{R} := \mathbb{Z}[X]/(X^M+1)$. The t-target* inhomogeneous short integer solution *problem relative to the NTRU trapdoor algorithm $\mathsf{TpdGen}_{\mathsf{NTRU}}$ with parameters $m,q,\alpha,\beta > 0$ is defined via the* t-$\mathcal{R}$-NTRU-ISIS *game from Fig. 6. We define the advantage of $\mathcal{A}$ in* t-$\mathcal{R}$-NTRU-ISIS *as*

$$\mathsf{Adv}^{\text{t-}\mathcal{R}\text{-NTRU-ISIS}}_{\mathcal{A},m,q,\alpha,\beta}(\lambda) := \Pr[1 \leftarrow \text{t-}\mathcal{R}\text{-NTRU-ISIS}_{\mathcal{A},m,q,\alpha,\beta}].$$

We may omit t, if it is 1.

Definition 21 (NTRU-SPISIS* [34, Slightly modified from Definition 10]). *Let $\mathcal{R} := \mathbb{Z}[X]/(X^M+1)$and let $t \geq 1$. The* second preimage inhomogeneous short integer solution *problem relative to the NTRU trapdoor algorithm $\mathsf{TpdGen}_{\mathsf{NTRU}}$ with parameters $q,\alpha,\beta > 0$ is defined via the* t-$\mathcal{R}$-NTRU-SPISIS* *game from Fig. 6. We define the advantage of $\mathcal{A}$ in* t-$\mathcal{R}$-NTRU-SPISIS* *as*

$$\mathsf{Adv}^{\text{t-}\mathcal{R}\text{-NTRU-SPISIS}^*}_{\mathcal{A},q,\alpha,\beta}(\lambda) := \Pr[1 \leftarrow \text{t-}\mathcal{R}\text{-NTRU-SPISIS}^*{}_{\mathcal{A},q,\alpha,\beta}].$$

Remark 2 (Relation to [34, Definition 10]). Our formulation differs from [34, Definition 10] in two respects. First, the syndromes are not uniform: we sample $\boldsymbol{u}_i, \boldsymbol{v}_i \leftarrow \mathcal{D}_{\mathcal{R},s}$ and set $\boldsymbol{c}_i := \boldsymbol{h}\boldsymbol{u}_i + \boldsymbol{v}_i$, so that $\boldsymbol{c}_i$ follows $\mathcal{Q}_{\boldsymbol{h}}$ and, conditioned on $\boldsymbol{c}_i$, the pair $(\boldsymbol{u}_i, \boldsymbol{v}_i)$ follows the coset Gaussian. These are exactly the distributions against which r_u and r_p are measured. Second, the game performs no rejection: only the returned preimage must be short. Rejection is done by the reduction instead, which reprograms every preimage sampling attempt and discards attempt i exactly when $\|(\boldsymbol{u}_i, \boldsymbol{v}_i)\|_2 > \beta$, mirroring the loop of Sign^+.

A divergence hop requires unconditioned samples: conditioning them costs up to $p^{-\alpha/(\alpha-1)}$ each. In [34] the conditioning is therefore never inside a hop. It is either a stopping rule applied identically on both sides of the hop, as in their $\mathsf{Game}_3 \to \mathsf{Game}_4$ and in their Theorem 1, where the random oracle is likewise programmed on every iteration, or it is introduced statistically, as in the transition preceding their reduction. Doing something similar in the QROM would place a conditioned distribution inside the reprogramming hop which is not possible with how we use the Rényi divergence.

The two assumptions are nevertheless close: on the targets whose given preimage happens to be short they differ only in the syndrome being drawn from $\mathcal{Q}_{\boldsymbol{h}}$ rather than $\mathcal{U}(\mathcal{R}_q)$, and the remaining targets have a large preimage as the only help. We therefore expect the hardness discussion of [34, Section 5.1] to apply here as well.

B Omitted Proofs

B.1 Proof of Theorem 3

Theorem 3. *An algorithm $\mathcal{A}$ making q quantum queries to O_{f_P} or O_{f_Q} has the same output distribution, up to statistical distance $8q\,\mathrm{dh}(P,Q) \le 8q\sqrt{\epsilon}$.*

We rely on techniques and perspectives on the QROM from [20] utilizing the compressed oracle framework from [66]. We recall notation from [20], simplified in [29]. Refer to the original work for a more detailed introduction.

Consider the multiregister $D = (D_x)_{x\in\mathcal{X}}$, where the state space of D_x is given by $\mathcal{H}_{D_x} = \mathbb{C}[\{0,1\}^n \cup \{\perp\}]$, meaning that it is spanned by an orthonormal set of vectors $|y\rangle$ labelled by $y \in \{0,1\}^n \cup \{\perp\}$. The initial state is set to be $|\perp\rangle := \bigotimes_x |\perp\rangle_{D_x}$. Let P be a probability distribution over $\{0,1\}^n$. Define the state $|\phi_P\rangle := \sum_{y\in\{0,1\}^n} \sqrt{P(y)}\,|y\rangle$ and the unitary

$$F^P = |\perp\rangle\langle\phi_P| + |\phi_P\rangle\langle\perp| + I_{\{0,1\}^n} - |\perp\rangle\langle\perp| - |\phi_P\rangle\langle\phi_P|\,.$$

When the compressed oracle is queried, a unitary O^P_{XYD}, acting on the query registers X and Y and the oracle register D, is applied, given by

$$O^P_{XYD} = \sum_{x\in\mathcal{X}} |x\rangle\langle x|_X \otimes O^{P,x}_{YD_x}, \quad \text{with} \quad O^{P,x}_{YD_x} = F^P_{D_x}\,\mathrm{CNOT}_{YD_x}\,F^P_{D_x}$$

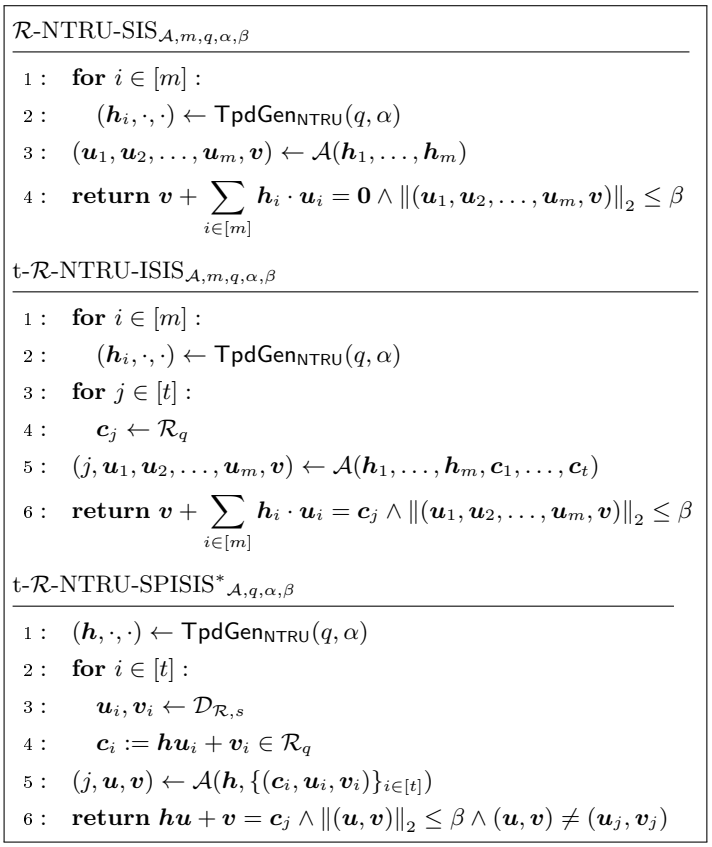

Fig. 6. Games defining $\mathcal{R}$-NTRU-SIS$_{\mathcal{A},m,q,\alpha,\beta}$, t-$\mathcal{R}$-NTRU-ISIS$_{\mathcal{A},m,q,\alpha,\beta}$ and t-$\mathcal{R}$-NTRU-SPISIS$^*{}_{\mathcal{A},q,\alpha,\beta}$.

where $\mathrm{CNOT}_{YD_x} |y\rangle |y_x\rangle = |y \oplus y_x\rangle |y_x\rangle$ for $y, y_x \in \{0,1\}^n$ and acting as the identity on $|y\rangle |\perp\rangle$. Instantiated with the uniform distribution over $\{0,1\}^n$, the formalism describes the "compressed" oracle as formalized in [29].

A q-query adversary $\mathcal{A}$ querying the compressed oracle with underlying distribution P can be described as in Fig. 7 where the measurement describes the final output of $\mathcal{A}$. The adversary describes the strategy via a unitary U[12] on LXY and the eventual outcome is described by a measurement on the state

$$|\psi_P\rangle := \underbrace{\left(UO^P_{XYD}\right)^q U}_{:=V_P} |0\rangle_{LXY} |\perp\rangle_D .$$

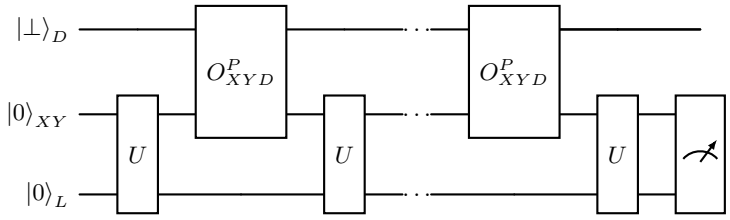

Fig. 7. General q-query quantum adversary using oracle with the underlying distribution of P.

Quantum Oracle Switching via Statistical Distance and the Compressed Oracle We are interested in the trace distance of $|\psi_P\rangle \langle\psi_P|$ and $|\psi_Q\rangle \langle\psi_Q|$ where we are given the statistical distance of P and Q.

Lemma 14. *Let ϵ be the statistical distance of P and Q over $\{0,1\}^n$. Then,*

$$\| |\psi_Q\rangle \langle\psi_Q| - |\psi_P\rangle \langle\psi_P| \|_{tr} \leq 8q \,\mathrm{dh}(P,Q) \leq 8q\sqrt{\epsilon}.$$

Proof. Both ψ_Q and ψ_P are constructed by applying V_Q or V_P to the same initial state. The trace norm can now be bounded by the operator norm of the unitaries that are applied to the initial state as in Eq. (7)

$$\| |\psi_Q\rangle \langle\psi_Q| - |\psi_P\rangle \langle\psi_P| \|_{tr} \leq 2\|V_Q - V_P\|_\infty .$$

By adding $2q$ null terms where we replace one F_Q with F_P, we get that the operator norm is bounded by the sum of the operator norms of the compressed oracle switches.

[12] One can also consider a sequence of unitaries U_i, but they can be combined into a single unitary with a counter and a controlled unitary to get the same behavior.

Let $j \in [2q]$ be the oracle switch, we are replacing in the current step. Now, add the null-term of the form

$$\pm \left(UO^P_{XYD}\right)^{q-\lfloor (j+1)/2 \rfloor} UW_j \left(UO^Q_{XYD}\right)^{\lfloor (j-1)/2 \rfloor} U$$

where

$$W_j := \begin{cases} \sum_x |x\rangle\langle x| \otimes F^P_{D_x}\, \mathrm{CNOT}_{YD_x}\, F^Q_{D_x} & \text{if } j \equiv 1 \mod 2, \\ \sum_x |x\rangle\langle x| \otimes F^Q_{D_x}\, \mathrm{CNOT}_{YD_x}\, F^Q_{D_x} & \text{if } j \equiv 0 \mod 2. \end{cases}$$

Now, we consider the original operator norm and add this nullterm for every $j \in [2q]$. By using the triangular inequality and the fact that the operator norm is invariant under application of unitaries, it follows that

$$\|V_Q - V_P\|_\infty \leq 2q \cdot \left\| \sum_x |x\rangle\langle x| \otimes (F^Q_{D_x} - F^P_{D_x}) \right\|_\infty .$$

We bound it in the next lemma. □

Lemma 15. *Let ϵ be the statistical distance of P and Q over $\{0,1\}^n$. Then,*

$$\left\| \sum_x |x\rangle\langle x| \otimes (F^Q_{D_x} - F^P_{D_x}) \right\|_\infty \leq 2\,\mathrm{dh}(P,Q) \leq 2\sqrt{\epsilon}.$$

Proof.

$$\begin{aligned} &\left\| \sum_x |x\rangle\langle x| \otimes (F^Q_{D_x} - F^P_{D_x}) \right\|_\infty = \max_x \left\| F^Q_{D_x} - F^P_{D_x} \right\|_\infty \\ &= \| |\perp\rangle\langle\phi_Q| - |\perp\rangle\langle\phi_P| + |\phi_Q\rangle\langle\perp| - |\phi_P\rangle\langle\perp| + |\phi_P\rangle\langle\phi_P| - |\phi_Q\rangle\langle\phi_Q| \|_\infty \\ &= \| |\perp\rangle (\langle\phi_Q| - \langle\phi_P|) + (|\phi_Q\rangle - |\phi_P\rangle)\langle\perp| + |\phi_P\rangle\langle\phi_P| - |\phi_Q\rangle\langle\phi_Q| \|_\infty \end{aligned}$$

We set $|\Delta_+\rangle := |\phi_Q\rangle + |\phi_P\rangle$ and $|\Delta_-\rangle := |\phi_Q\rangle - |\phi_P\rangle$. Now we express the above using these orthogonal states.

$$\begin{aligned} \left\| \sum_x |x\rangle\langle x| \otimes (F^Q_{D_x} - F^P_{D_x}) \right\|_\infty = \Big\| & |\perp\rangle\langle\Delta_-| + |\Delta_-\rangle\langle\perp| \\ & + 1/4(|\Delta_+\rangle - |\Delta_-\rangle)(\langle\Delta_+| - \langle\Delta_-|) \\ & - 1/4(|\Delta_+\rangle + |\Delta_-\rangle)(\langle\Delta_+| + \langle\Delta_-|) \Big\|_\infty \end{aligned}$$

Considering an orthonormal basis that includes normalized versions of $|\perp\rangle, |\Delta_-\rangle, |\Delta_+\rangle$, it is sufficient to consider how the operator behaves:

$$\begin{aligned}
|\perp\rangle &\mapsto |\Delta_-\rangle \\
\frac{1}{\|\Delta_-\|_2}|\Delta_-\rangle &\mapsto \frac{1}{\|\Delta_-\|_2}\left(\langle\Delta_-|\Delta_-\rangle|\perp\rangle - 1/2\,\langle\Delta_-|\Delta_-\rangle|\Delta_+\rangle\right) \\
&= \|\Delta_-\|_2\left(|\perp\rangle - 1/2\,|\Delta_+\rangle\right) \\
\frac{1}{\|\Delta_+\|_2}|\Delta_+\rangle &\mapsto \frac{1}{\|\Delta_+\|_2}\left(-1/2\,\langle\Delta_+|\Delta_+\rangle|\Delta_-\rangle\right) = -1/2\|\Delta_+\|_2\,|\Delta_-\rangle
\end{aligned}$$

In matrix form, we thus get

$$M = \begin{pmatrix} 0 & \|\Delta_-\|_2 & 0 \\ \|\Delta_-\|_2 & 0 & -\frac{1}{2}\|\Delta_-\|_2\|\Delta_+\|_2 \\ 0 & -\frac{1}{2}\|\Delta_-\|_2\|\Delta_+\|_2 & 0 \end{pmatrix}.$$

The operator $F^Q_{D_x} - F^P_{D_x}$ is Hermitian, so M is (real) symmetric, and its operator norm is the largest absolute value among its eigenvalues. Writing $a := \|\Delta_-\|_2$ and $b := \frac{1}{2}\|\Delta_-\|_2\|\Delta_+\|_2$, the characteristic polynomial of M is $-\lambda(\lambda^2 - (a^2 + b^2))$, so its eigenvalues are 0 and $\pm\sqrt{a^2+b^2}$, giving $\|M\|_\infty = \sqrt{a^2+b^2} = \|\Delta_-\|_2\sqrt{1+\frac{1}{4}\|\Delta_+\|_2^2}$. Trivially $\|\Delta_+\|_2 \leq 2$, and thus $\|M\|_\infty \leq \sqrt{2}\|\Delta_-\|_2$.

We thus have

$$\left\|\sum_x |x\rangle\langle x| \otimes (F^Q_{D_x} - F^P_{D_x})\right\|_\infty \leq \sqrt{2}\||\phi_P\rangle - |\phi_Q\rangle\|_2.$$

Now we can proceed considering that both are unit vectors and we get

$$\||\phi_P\rangle - |\phi_Q\rangle\|_2 = \sqrt{2 - \langle\phi_P|\phi_Q\rangle - \langle\phi_Q|\phi_P\rangle}.$$

With $\langle\phi_P|\phi_Q\rangle = \langle\phi_Q|\phi_P\rangle$, we can bound this term separately. As $\langle\phi_P|\phi_Q\rangle = \sum_{y\in\{0,1\}^n}\sqrt{P(y)Q(y)} = \mathrm{BC}(P,Q)$ we get that

$$\begin{aligned}
\sqrt{2 - \langle\phi_P|\phi_Q\rangle - \langle\phi_Q|\phi_P\rangle} &= \sqrt{2}\sqrt{1 - \mathrm{BC}(P,Q)} \\
&= \sqrt{2}\,\mathrm{dh}(P,Q) \leq \sqrt{2\epsilon},
\end{aligned}$$

and in total

$$\left\|\sum_x |x\rangle\langle x| \otimes (F^Q_{D_x} - F^P_{D_x})\right\|_\infty \leq 2\,\mathrm{dh}(P,Q) \leq 2\sqrt{\epsilon}.$$

□

This makes it possible to get to our main result.

Theorem 11. *Let $\mathcal{A}$ be a quantum algorithm given oracle access to an oracle instantiated with P or Q and*

$$|\Pr[\mathcal{A} \to 1|P] - \Pr[\mathcal{A} \to 1|Q]| \geq \zeta,$$

where $\mathcal{A}$ can make q queries. Then $\Delta(P,Q) = \Omega(\zeta^2/q^2)$.

Proof. View the algorithm in the compressed oracle formalism. This perfectly simulates the interaction with the oracle. The final output of $\mathcal{A}$ corresponds to a measurement on the final state, which by Lemma 14 has trace distance $8q\sqrt{\epsilon}$ where ϵ is the statistical distance of P and Q. The trace distance also provides an upper bound for the difference in output probability, such that

$$|\Pr[\mathcal{A} \to 1|P] - \Pr[\mathcal{A} \to 1|Q]| \leq 8q\sqrt{\epsilon}.$$

Rewriting it, we get

$$\zeta \leq 8q\sqrt{\epsilon} \Rightarrow \zeta^2/(64q^2) \leq \epsilon = \Delta(P,Q).$$

□

This result improves the bound from [64, Section 7.2] by a factor of q, which itself improved the bound by a factor of q compared to [12]. After finishing this, we found that there is another work achieving the same asymptotic bounds [7], but using the quantum query complexity. Our approach therefore only gives a different proof technique to reach the same final statement, albeit they do not provide explicit bounds needed in cryptography. Their result, however, only provides a bound for constant advantage, and our result gives a bound for any advantage ζ.

B.2 Proof of Proposition 1

The counterexample presented in Section 3.2 shows that while the ratio of success probabilities increases, the absolute success probability decreases. This opens the door to an approach with a mix of additive and multiplicative loss terms.

We use the small-range distributions and subsequent results from [64] towards such a mixed approach. We note that this approach opens a path to exploiting a small Rényi divergence even between distributions that have large statistical distance. The small-range distribution approach replaces the i.i.d. function with the composition of two i.i.d functions where the first one is uniform with small range r, and the second one uses the actual distribution. This introduces a loss based on $\mathcal{A}$'s capabilities of distinguishing this oracle from the honest oracle. The small-range distribution with range size r gets r classical samples from the original i.i.d. function, this approach is amenable to applying the classical Rényi divergence machinery.

Formally, given a distribution D on $\mathcal{Y}$, define $\mathsf{SR}_r^D(\mathcal{X})$ as the following distribution on functions from $\mathcal{X}$ to $\mathcal{Y}$:

– For each $i \in [r]$, sample a value $y_i \in \mathcal{Y}$ according to D.
– For each $x \in \mathcal{X}$, pick a uniformly random $i \in [r]$ and set $O(x) = y_i$.

Lemma 16 ([64, Corollary 7.5]). *The output distributions of a quantum algorithm making q quantum queries to an oracle either drawn from SR_r^D or $D^{\mathcal{X}}$[13] have statistical distance at most $\ell(q)/r$, where $\ell(q) = \pi^2(2q)^3/3 < 27q^3$.*

Proposition 1. *Let $\alpha \in (1, \infty], r \in \mathbb{N}_{\geq 1}$ and P, Q be classical distributions over $\mathcal{Y}$ with $\mathrm{Supp}(P) \subseteq \mathrm{Supp}(Q)$. For any $\mathcal{A}$ making at most q quantum queries*

$$\Pr\left[1 \leftarrow \mathcal{A}^{f_P}\right] \leq \left(R_\alpha(P\|Q)^r \left(\Pr\left[1 \leftarrow \mathcal{A}^{f_Q}\right] + \frac{\ell(q)}{r}\right)\right)^{\frac{\alpha-1}{\alpha}} + \frac{\ell(q)}{r},$$

where $\ell(q) = \pi^2(2q)^3/3 < 27q^3$.

Proof (Proof of Proposition 1). Apply Lemma 16, then apply post-processing of the Rényi divergence on the classical samples of which there are r, and then apply Lemma 16 again:

$$\begin{aligned}
\Pr\left[1 \leftarrow \mathcal{A}^{f_P}\right] &\leq \Pr\left[1 \leftarrow \mathcal{A}^{\mathsf{SR}_r^P}\right] + \frac{\ell(q)}{r} \\
&\leq \left(\delta^r \Pr\left[1 \leftarrow \mathcal{A}^{\mathsf{SR}_r^Q}\right]\right)^{\frac{\alpha-1}{\alpha}} + \frac{\ell(q)}{r} \\
&\leq \left(\delta^r \left(\Pr\left[1 \leftarrow \mathcal{A}^{f_Q}\right] + \ell(q)/r\right)\right)^{\frac{\alpha-1}{\alpha}} + \frac{\ell(q)}{r}. \qquad (10)
\end{aligned}$$

□

Applicability This technique can be used in some regimes, but in some situations, using a generic relationship between Rényi divergence and the statistical distance in combination with Theorem 3 will yield stronger bounds. Consider a Rényi divergence of factor $(1+\epsilon)$ which is applied r times. The multiplicative factor becomes

$$(1+\epsilon)^r = e^{r\ln(1+\epsilon)} \approx e^{r\epsilon}.$$

The factor should not be super-polynomial in q, so ϵ must be chosen as $\frac{k \ln q}{r}$ where k is some positive constant. The factor $\ell(q)/r$ should be negligible, forcing a super-polynomial r. Therefore, ϵ needs to be selected to be negligible. Following standard inequalities of the Rényi divergence (see Appendix A.1), it bounds the statistical distance as $\Delta(P, Q) \leq \sqrt{\ln(1+\epsilon)} \approx \sqrt{\epsilon}$, which is negligible.

Using Pinsker's inequality and Theorem 3 we can obtain a bound on the trace distance, namely

$$C \cdot q \cdot \sqrt{\sqrt{\epsilon}},$$

for some small constant C. In Eq. (10), clearly, $r \leq -\log \Pr\left[1 \leftarrow \mathcal{A}^{f_Q}\right]/\epsilon$, otherwise the $\delta^r \Pr\left[1 \leftarrow \mathcal{A}^{f_Q}\right] \geq 1$. This implies that the bound in Lemma 16 is at least

[13] This notation corresponds to our oracle notation O_{f_D}.

$$\frac{\ell(q)}{r} \geq \frac{q^3}{-\log \Pr\left[1 \leftarrow \mathcal{A}^{Q^{\mathcal{X}}}\right]/\epsilon} = \epsilon \frac{q^3}{-\log \Pr[1 \leftarrow \mathcal{A}^{f_Q}]}.$$

The statistical distance error bound $C \cdot q \cdot \sqrt{\sqrt{\epsilon}}$ is smaller than the additive error of the small range distribution technique at least whenever

$$q \geq \sqrt{\frac{-C \log(\Pr[1 \leftarrow \mathcal{A}^{f_Q}])}{\epsilon^{3/4}}}.$$

B.3 Proof of Lemma 6

Lemma 6 (Adaptive Reprogramming with Distribution Switching). *Let $\mathcal{X}_1, \mathcal{X}_2, \mathcal{X}'$ and $\mathcal{Y}$ be finite sets, and let p be a distribution on $\mathcal{X}_1 \times \mathcal{X}'$. Let Q be a distribution over $\mathcal{Y}$ and Dist be any distinguisher, issuing q (quantum) queries to O and R reprogramming instructions such that each instruction consists of a value x_2, together with the fixed distribution p. Then*

$$\left|\Pr\left[1 \leftarrow \mathsf{REPRO}_0^{\mathsf{Dist}}\right] - \Pr\left[1 \leftarrow \mathsf{REPRO}_{1,Q}^{\mathsf{Dist}}\right]\right| \leq R \cdot \Delta(Q, \mathcal{U}(\mathcal{Y})) + \delta_{repr}$$
$$\Pr\left[1 \leftarrow \mathsf{REPRO}_0^{\mathsf{Dist}}\right] \leq \left(R_\alpha(\mathcal{U}(\mathcal{Y}) \| Q)^R \Pr\left[1 \leftarrow \mathsf{REPRO}_{1,Q}^{\mathsf{Dist}}\right]\right)^{\frac{\alpha-1}{\alpha}} + \delta_{repr}$$

where $\delta_{repr} = \frac{3R}{2}\sqrt{q \cdot p_{\max}}$ with $p_{\max} := \max_{x_1} p(x_1)$ and $\alpha \in (1, \infty]$.

Proof (Proof of Lemma 6). The proof is quite straightforward. Apply Lemma 4 and sample all values that need resampling. This gives the additive terms from the reprogramming. Replace the R classical samples from the uniform distribution, by values from distribution Q and use the properties of the Rényi divergence. The factor of $R_\alpha(\mathcal{U}(\mathcal{Y})\|Q)^R$ follows from observing that each sample to be reprogrammed is independent, so we can apply Lemma 2 to get the overall bound on the Rényi divergence.

$$\begin{aligned}
\left|\Pr\left[1 \leftarrow \mathsf{REPRO}_{1,Q}^{\mathsf{D}}\right] - \Pr\left[1 \leftarrow \mathsf{REPRO}_0^{\mathsf{D}}\right]\right| & \\
&\leq \left|\Pr\left[1 \leftarrow \mathsf{REPRO}_1^{\mathsf{D}}\right] - \Pr\left[1 \leftarrow \mathsf{REPRO}_0^{\mathsf{D}}\right]\right| \\
&\quad + \left|\Pr\left[1 \leftarrow \mathsf{REPRO}_{1,Q}^{\mathsf{D}}\right] - \Pr\left[1 \leftarrow \mathsf{REPRO}_1^{\mathsf{D}}\right]\right| \\
&\leq \delta_{repr} + R\Delta(Q, \mathcal{U}(\mathcal{Y})) \\
\Pr\left[1 \leftarrow \mathsf{REPRO}_0^{\mathsf{D}}\right] &\leq \Pr\left[1 \leftarrow \mathsf{REPRO}_1^{\mathsf{D}}\right] + \delta_{repr} \\
&\leq \left(R_\alpha(\mathcal{U}(\mathcal{Y})\|Q)^R \Pr\left[1 \leftarrow \mathsf{REPRO}_{1,Q}^{\mathsf{Dist}}\right]\right)^{\frac{\alpha-1}{\alpha}} + \delta_{repr}
\end{aligned}$$

where $\delta_{repr} = \frac{3R}{2}\sqrt{qp_{\max}}$. □

B.4 Proof of Theorem 8

Theorem 8. *Let Σ be a special-sound Σ-protocol with HVZK simulator* Sim *and simulator commitment min-entropy*[14] *$\beta(\lambda)$. Consider the AOS RSS with maximal ring size $\kappa = \kappa(\lambda) \geq N$. For any* SUF-CRA *adversary $\mathcal{A}$ making at most q_s signing queries and* ANON *adversary $\mathcal{B}$ making at most q_c challenge queries, each making at most q_{H} quantum queries to the RO* H*, there exists a* UF-NRA *adversary $\mathcal{A}_{\text{UF-NRA}}$, a* CUR *adversary $\mathcal{A}_{\text{CUR}}$ and a* W-REC *adversary $\mathcal{A}_{\text{W-REC}}$ such that $\mathsf{Adv}^{\text{SUF-CRA}}_{\mathcal{A},\mathsf{AOS}(\Sigma),N,q_s}(\lambda)$ and $\mathsf{Adv}^{\text{ANON}}_{\mathcal{B},\mathsf{AOS}(\Sigma),N,q_c}(\lambda)$ are bounded as in Table 5 and the parameters are defined as in Table 6.*

Proof. For SUF-CRA security, we define a sequence of games. The proof follows essentially the same pattern as [38, Theorem 3].

The first game is the normal SUF-CRA game. In the second game, we make two changes. The jth signature query $\mathsf{OSign}(i_j, m_j, \rho_j)$ will be modified as follows. We sample a value $\mathsf{ch}_j \leftarrow \mathcal{C}$ uniformly at random and everything in the ring stays as it is, until the query to $\mathsf{H}(i_j, \rho_j, \mathsf{com}_{i_j-1}, m)$ is made. This query will now be programmed to the initially sampled value of ch_j. All values ch_j for $j \in [q_s]$ can be sampled at the beginning. In the next game, we have already defined ch_j at the beginning, so we can use the HVZK simulator for every participant of the ring and get a valid signature. This removes the secret key, and the oracle can be simulated, i.e., we can simulate the oracle and transition to UF-NRA. The games are formally depicted in Fig. 8.

Game_0: This is the original game, so here we have

$$\mathsf{Adv}^{\text{SUF-CRA}}_{\mathcal{A},\mathsf{AOS}(\Sigma),N,q_s}(\lambda) = \Pr\left[\mathsf{Game}_0^{\mathcal{A}}\right].$$

Game_1: In game Game_1, each signing query is adapted as described in Fig. 8. The adaptive reprogramming lemma, Lemma 4 from [38], yields a bound on the difference of Game_0 and Game_1. At the cost of at most N additional queries, assume that the adversary runs the verifier on its output signature. We construct a reprogramming distinguisher: Run $\mathcal{A}$, calling the reprogramming oracle for each signing query. The distinguisher issues $q_H + \kappa \cdot q_s + N$ queries to H with q_s many reprogramming queries, so Lemma 4 yields

$$|\Pr[\mathsf{Game}_0(\mathcal{A})] - \Pr[\mathsf{Game}_1(\mathcal{A})]| \leq \frac{3q_s}{2}\sqrt{(q_{\mathsf{H}} + \kappa \cdot q_s + N) \cdot p_{\max}} + \epsilon'_{repr}$$

where $p_{\max} \leq 2^{-\beta(\lambda)}$ by assumption on the simulator commitment min-entropy, and ϵ'_{repr} is an additional error as the eventually produced signature forgery must be valid for the non-reprogrammed hash function H. For the eventual signature (ρ^*, m^*, σ^*) to be valid, there can not be a pair (ρ^*, m^*, σ') that was produced by an oracle query where i was the signer and $\mathsf{com}'_{i-1} = \mathsf{com}^*_{i-1}$. For this index,

[14] If the min-entropy is too small, we can either add commitment entropy by appending a random string, or we can instantiate the AOS construction with a *salted* RO, resulting in a higher effective commitment min-entropy.

Unforgeability games $\mathsf{Game}_0 - \mathsf{Game}_2$

$q \leftarrow 0; \mathcal{L} \leftarrow \emptyset$
for $i \in [N]$:
  $(\mathsf{inst}_i, w_i) \leftarrow \mathsf{Gen}(1^\lambda)$
$\rho = \{\mathsf{inst}_1, \ldots, \mathsf{inst}_N\}$
$(\rho^*, m^*, \sigma^*) \leftarrow \mathcal{A}^{\mathsf{OSign}}(\rho)$
if $(\rho^*, m^*, \sigma^*) \in \mathcal{L} \vee \rho^* \not\subseteq \rho$:
  return false
$\{\mathsf{inst}_1^*, \ldots, \mathsf{inst}_{N^*}^*\} \leftarrow \rho^*$
$((\mathsf{com}_i^*, \mathsf{rsp}_i^*))_{i\in[N^*]} \leftarrow \sigma^*$
for $i \in [N^*]$:
  $\mathsf{ch}_{i+1}^* \leftarrow \mathsf{H}(i+1, \rho^*, \mathsf{com}_i^*, m^*)$
return $\bigwedge_{i\in[N^*]} \mathsf{V}_2(\mathsf{inst}_i^*, \mathsf{com}_i^*, \mathsf{ch}_i^*, \mathsf{rsp}_i^*)$

Anonymity games Game_0 to Game_2

$q \leftarrow 0; b \leftarrow \{0,1\}$
for $i \in [N]$:
  $(\mathsf{inst}_i, w_i) \leftarrow \mathsf{Gen}(1^\lambda)$
$\rho = \{\mathsf{inst}_i\}_{i\in[N]}; \mathsf{SK} = \{w_i\}_{i\in[N]}$
$b' \leftarrow \mathcal{B}^{\mathsf{CHAL}}(\rho, \mathsf{SK})$
return $b = b'$

$\mathsf{CHAL}(i_0, i_1, \rho' = \{\mathsf{inst}'_i\}_{i\in[N']}, m)$

if $\mathsf{inst}'_{i_0} \notin \rho \vee \mathsf{inst}'_{i_1} \notin \rho$:
  return $\perp$
return $\mathsf{OSign}(i_b, \rho', m)$

$\mathsf{OSign}(i, \rho' = \{\mathsf{inst}'_1, \ldots, \mathsf{inst}'_{N'}\}, m)$

if $q = q_{\max} \vee i \notin [N'] \vee \mathsf{inst}'_i \notin \rho$: // $q_{\max}$ is q_s or q_c in the respective games
  return $\perp$
$q \leftarrow q + 1$
// let w be the witness such that $\mu(w) = \mathsf{inst}'_i$
$\mathsf{ch}_i \leftarrow \mathcal{C}$ // $\mathsf{Game}_1 - \mathsf{Game}_2$
$(\mathsf{com}_i, \mathsf{st}) \leftarrow \mathsf{P}_1(\mathsf{inst}'_i)$ // $\mathsf{Game}_0 - \mathsf{Game}_1$
$(\mathsf{com}_i, \mathsf{rsp}_i) \leftarrow \mathsf{Sim}(\mathsf{inst}'_i, \mathsf{ch}_i)$ // Game_2
for $j = i+1, \ldots, N', 1, \ldots, i-1$:
  $\mathsf{ch}_j \leftarrow \mathsf{H}(j, \rho', \mathsf{com}_{j-1}, m)$
  $(\mathsf{com}_j, \mathsf{rsp}_j) \leftarrow \mathsf{Sim}(\mathsf{inst}'_j, \mathsf{ch}_j)$
$\mathsf{ch}_i \leftarrow \mathsf{H}(i, \rho', \mathsf{com}_{i-1}, m)$ // Game_0
$\mathsf{rsp}_i \leftarrow \mathsf{P}_2(\mathsf{state}, w, \mathsf{ch}_i)$ // $\mathsf{Game}_0 - \mathsf{Game}_1$
$\mathsf{H} := \mathsf{H}^{(i,\rho',\mathsf{com}_{i-1},m)\mapsto \mathsf{ch}_i}$ // $\mathsf{Game}_1 - \mathsf{Game}_2$
$\sigma \leftarrow ((\mathsf{com}_1, \mathsf{rsp}_1), \ldots, (\mathsf{com}_{N'}, \mathsf{rsp}_{N'}))$
$\mathcal{L} \leftarrow \mathcal{L} \cup \{(\rho', m, \sigma)\}$
return σ

Fig. 8. Unforgeability and anonymity games Game_0 to Game_2 in Theorem 8. All indices are interpreted to be modulo N, N', and modulo N^*, respectively. In Game_0, we have the normal unforgeability/anonymity game. In Game_1, the RO is reprogrammed to a RO value, and in Game_2, the signature is generated entirely without the secret key w.

there was reprogramming, so the actual signature might not be valid under the non-reprogrammed oracle. We bound this probability as ϵ'_{repr}. Assume a signing query exists where $\mathsf{com}'_{i-1} = \mathsf{com}^*_{i-1}$. Let ch^*_j and ch'_j for $j \in [N^*]$ be challenges computed with the reprogrammed oracle H. If $\mathsf{ch}'_{i-1} \neq \mathsf{ch}^*_{i-1}$ and both signatures are valid, we can extract the witness via special soundness, creating an adversary $\mathcal{A}_{\text{W-REC}}$. Now assume that $\mathsf{ch}'_{i-1} = \mathsf{ch}^*_{i-1}$, then, either the same commitments for the previous ring member were used, i.e., $\mathsf{com}'_{i-2} = \mathsf{com}^*_{i-2}$, or a hash collision was found. For a hash collision with ring ρ^*, there are at most $q_\mathsf{H} + N^* \cdot q_s + N^*$ queries to H and a collision happens with probability at most $20(q_\mathsf{H} + N \cdot q_s + N)^3/|\mathcal{C}|$ using Lemma 9.[15] If the commitments are the same for $i-2$, we can repeat the same argument. Assuming no hash collisions, all commitments and (challenges) are the same. If all commitments and challenges are the same, then there must be a response that is different for one index $j \in [N^*]$ because the signatures are different. This can be used to construct a CUR adversary $\mathcal{A}_{\text{CUR}}$ against Σ. The W-REC and CUR adversaries rely on guessing the correct index when placing the Σ-protocol instances before giving them to the adversary, which introduces the multiplicative factor of N. Altogether, this allows to bound

$$\epsilon'_{repr} \leq N \cdot (\mathsf{Adv}^{\text{CUR}}_{\mathcal{A}_{\text{CUR}},\Sigma}(\lambda) + \mathsf{Adv}^{\text{W-REC}}_{\mathcal{A}_{\text{W-REC}},\Sigma}(\lambda)) + 20(q_\mathsf{H} + N \cdot q_s + N)^3/|\mathcal{C}|.$$

Game_2: In game Game_2, we remove the signing key. We do not compute an initial commitment, but rather compute the commitment and response from the challenge using the simulator, and reprogram as in the previous game.

Lets consider the statistical distance approach first. There are q_s queries, and for each replacement, the statistical difference is at most $\epsilon^{\text{HVZK}}_\Sigma$ such that

$$|\Pr[\mathsf{Game}_1(\mathcal{A})] - \Pr[\mathsf{Game}_2(\mathcal{A})]| \leq q_s \cdot \epsilon^{\text{HVZK}}_\Sigma.$$

$\mathcal{A}$ now makes no actual signing queries, i.e., we can define a UF-NRA adversary $\mathcal{A}_{\text{UF-NRA}}$ that plays the role of the challenger for $\mathcal{A}$ by simulating all transcripts themselves and reprogramming the RO

$$\Pr[\mathsf{Game}_2(\mathcal{A})] \leq \mathsf{Adv}^{\text{UF-NRA}}_{\mathcal{A}_{\text{UF-NRA}},\mathsf{AOS}(\Sigma),N}(\lambda).$$

For the Rényi divergence approach we can use Lemmas 1 and 2 as the underlying distributions are independent. This introduces the multiplicative factor $(\delta^{\alpha\text{-HVZK}}_\Sigma)^{q_s}$.

We now prove anonymity. Classically, we would rely on reprogramming and running the simulator. We start with the normal anonymity game $\mathsf{Game}_0 = \text{ANON}_{\mathcal{B},\mathsf{RSig},N,q_c}(\lambda)$, and from it, we define two small modifications formally depicted in Fig. 8. The first one is just a reprogramming of the QROM, and the second game uses the simulator to remove the use of the secret key. With

[15] Here, we make use of domain-separation between different rings. The hash function we are using in the definition maps directly into $\mathcal{C}$. Note that γ can be chosen such that hashing into $\mathcal{Y}$ and mapping to $\mathcal{C}$ with γ satisfies $\gamma_{cl}/\mathcal{Y} \leq \frac{2}{|\mathcal{C}|}$.

the same reasoning as in the unforgeability part, except for the additional hash assumption that was needed for the eventual forgery, it follows that

$$|\Pr[\mathsf{Game}_0(\mathcal{B})] - \Pr[\mathsf{Game}_1(\mathcal{B})]| \leq 3q_c/2\sqrt{(q_{\mathsf{H}} + \kappa \cdot q_c) \cdot 2^{-\beta(\lambda)}}.$$

We distinguish two cases now. First, we can use the statistical distance, and observe that the distance between a signature from the oracle in the two games is the difference of using the simulator or the actual proof algorithm for one of the ring members. Hence, the statistical distance for each oracle query is bounded by $\epsilon_{\Sigma}^{\text{HVZK}}$. There are q_c queries which yield the final bound of $q_c \cdot \epsilon_{\Sigma}^{\text{HVZK}}$. Second, we can use the Kullback-Leibler divergence to bound the difference in probability between Game_1 and Game_2 using Lemma 3. This introduces a factor of $\sqrt{\left(q_c \cdot \epsilon_{\Sigma}^{KL\text{-HVZK}}\right)/2}$. Considering Game_2, the signing oracle is independent of the choice of b, so the probability that $\mathcal{B}$ outputs b is exactly $1/2$. □

B.5 Proof of Theorem 9

Theorem 9 (UF-NRA using Measure-and-Reprogram 2.0). *Let Σ be a Σ-protocol. For any* UF-NRA *adversary $\mathcal{A}$ making at most q_{H} quantum queries to H, there exists an impersonation adversary $\mathcal{B}$ that can convince an honest verifier V such that*

$$\mathsf{Adv}^{\text{UF-NRA}}_{\mathcal{A},\mathsf{AOS}(\Sigma),N}(\lambda) \leq N \cdot (2(q_{\mathsf{H}} + N) + 1)^{2N} \cdot \mathsf{Adv}^{\text{IMP}}_{\mathcal{B},\Sigma}(\lambda).$$

Proof. We construct an adversary $\mathcal{B}$ from $\mathcal{A}$. $\mathcal{B}$ receives an instance inst, where $(\mathsf{inst}, w) \leftarrow \Sigma.\mathsf{Gen}(1^{\lambda})$.

$\mathcal{B}$ samples $i \leftarrow [N]$, sets $\mathsf{inst}_i = \mathsf{inst}$ and generates $(\mathsf{inst}_j, w_j) \leftarrow \Sigma.\mathsf{Gen}(1^{\lambda})$ for $j \in [N] \setminus \{i\}$. It forwards $(\mathsf{inst}_k)_{k\in[N]}$ to $\mathcal{A}$. Since all instances are identically distributed, $\mathcal{A}$'s view is independent of i.

Let ρ^* with $N' = |\rho^*| \leq N$ be the ring of $\mathcal{A}$'s final forgery. We ensure that $\mathcal{A}$ must have made at least a classical query for every $(1, \rho^*, \mathsf{com}^*_{N'}, m^*), \ldots, (N', \rho^*, \mathsf{com}^*_{N'-1}, m^*)$ from the final forgery. Theorem 1 yields an order π of the queries in which we may reprogram, at a multiplicative loss of $(2(q_{\mathsf{H}} + N) + 1)^{2N}$. This mimics the time-ordered list that is immediate in the classical case. One reprogramming carries the Σ-protocol challenge; the remaining $N-1$ only commit $\mathcal{A}$ to a classical ordering.

Recall that the query with index j takes com_{j-1} as input and defines ch_j as its output. Hence, for $\mathcal{B}$ to place its instance at a position p of ρ^*, it must read com_p from the input of query $p+1$ and program the challenge received from V as the output of query p, which requires $\pi(p+1) < \pi(p)$ (indices modulo N'). Such a p always exists. As π is independent of i, the guess i lands on a position that lies in ρ^* and satisfies this condition with probability at least $1/N$.

For every other query j that must be reprogrammed, $\mathcal{B}$ samples $\mathsf{ch}_j \leftarrow \mathcal{C}$ uniformly and programs it into the oracle. When $\mathcal{A}$ outputs a forgery $((\mathsf{com}_1, \mathsf{rsp}_1), \ldots, (\mathsf{com}_{N'}, \mathsf{rsp}_{N'}))$, $\mathcal{B}$ sends rsp_p to V. If the signature verifies, then in particular $\mathsf{V}_2(\mathsf{inst}_p, \mathsf{com}_p, \mathsf{ch}, \mathsf{rsp}_p)$ holds, which concludes the proof. □

B.6 Proof of Lemma 8

Lemma 8. *Let* $\gamma_{cl} := \max_{c \in \mathcal{C}} |\{y \in \mathcal{Y} \mid \gamma(y) = c\}|$, *then*

$$[\![\mathsf{SZ}_{\leq s} \setminus \mathsf{CL} \to \mathsf{CL}]\!] \leq \sqrt{10}\sqrt{s \cdot \gamma_{cl}/|\mathcal{Y}|}.$$

Proof (Lemma 8). We apply Theorem 2 with $\mathsf{P} := \mathsf{SZ}_{\leq s} \setminus \mathsf{CL}$ and $\mathsf{P}' := \mathsf{CL}$. Consider an arbitrary $D \in \mathcal{D}$ and input $x \in \mathcal{X}$. By [28, Remark 2.5], the argument of the maximum in the RHS of Eq. (1) does not depend on the value $D(x)$, but only on the values of D outside x. Therefore, we may assume that $D(x) = \bot$ when bounding $\Pr\left[U \in L^{x,D}\right]$. Since $\mathsf{P} \subset \mathsf{SZ}_{\leq s}$, if $\mathsf{P}|_{D|^x}$ is non-empty, it follows that $D \in \mathsf{SZ}_{\leq s}$. We thus assume $D \in \mathsf{SZ}_{\leq s}$, otherwise, Theorem 2 allows us to set $L^{x,D} = P|_{D|^x} = \emptyset$. The convention of $D(x) = \bot$ implies that $D \in \mathsf{CL} \Rightarrow D[x \mapsto U] \in \mathsf{CL}$ for all $U \in \mathcal{Y}$. Hence, with the same argument as before, we may assume that $D \notin \mathsf{CL}$, as we could set $L^{x,D} = P|_{D|^x}$ which is empty.

It remains to consider $D \in \neg\mathsf{CL} \cap \mathsf{SZ}_{\leq s}$. Assuming $D(x) = \bot$, we have that $\bot \notin P'|_{D|^x}$, and so we choose $L^{x,D} = P'|_{D|^x}$. We observe that

$$U \in L^{x,D} \Leftrightarrow D[x \mapsto U] \in P' \Rightarrow \exists x' : D(x') \neq \bot \wedge \gamma \circ D(x') = \gamma \circ D(x).$$

As $D \in \mathsf{SZ}_{\leq s}$, there are at most s possible values x' with $D(x') \neq \bot$. Therefore, we consider $\Pr[\gamma(U) \in \{\gamma(D(x')) \mid x' : D(x') \neq \bot\}]$. This holds if and only if $U \in \gamma^{-1}(\{\gamma(D(x')) \mid x' : D(x') \neq \bot\})$ which is of size at most $s \cdot \gamma_{cl}$. Thus,

$$\Pr\left[U \in L^{x,D}\right] \leq s \cdot \gamma_{cl}/|\mathcal{Y}|$$

holds. Applying Theorem 2, this yields

$$[\![\mathsf{SZ}_{\leq s} \setminus \mathsf{CL} \to \mathsf{CL}]\!] \leq \sqrt{10}\sqrt{s \cdot \gamma_{cl}/|\mathcal{Y}|}.$$

□

B.7 Proof of Lemma 10

Lemma 10. *Let* $\gamma_{cl} := \max_{c \in \mathcal{C}} |\{y \in \mathcal{Y} \mid \gamma(y) = c\}|$, *then*

$$[\![\mathsf{SZ}_{\leq s} \setminus \mathsf{SUC} \setminus \mathsf{CL} \to \mathsf{SUC}]\!] \leq \sqrt{10}\sqrt{\max\left\{N \cdot \ell, s \cdot \ell, \gamma_{cl} \cdot (sz^{\mathfrak{S}}_{triv})^N\right\}/|\mathcal{Y}|}.$$

Proof. We apply Theorem 2 with $\mathsf{P} := \mathsf{SZ}_{\leq s} \setminus \mathsf{SUC} \setminus \mathsf{CL}$ and $\mathsf{P}' := \mathsf{SUC}$. Consider an arbitrary but fixed $D \in \mathcal{D}$ and input $x \in \mathcal{X}$. Furthermore, by [28, Remark 2.5], the transition capacity does not depend on the value $D(x)$, but only on the values of D outside x. Therefore, we may assume that $D(x) = \bot$ for any x for which we consider the transition capacity. When considering P, we observe that it satisfies the database property $\mathsf{SZ}_{\leq s}$. If $\mathsf{P}|_{D|^x}$ is non-empty, the database D must satisfy $\mathsf{SZ}_{\leq s}$, i.e., it is bounded in size.

We now consider different cases as depicted in Table 7. We can assume $D \in \mathsf{SZ}_{\leq s}$. Otherwise, Theorem 2 allows us to set $L^{x,D} = P|_{D|^x}$ which is empty. The

Table 7. We have to consider two types of transitions regarding success, namely the transitions from SUC to $\neg\mathsf{SUC}$ and the transitions from $\neg\mathsf{SUC}$ to SUC. We separate the individually domain-separated commitment and challenge queries.

Case	Condition	Query Type	Bound
$1a$	$D \in \mathsf{SUC} \cap \neg\mathsf{CL} \cap \mathsf{SZ}_{\leq s}$	commitment	$\leq N \cdot \ell/\lvert\mathcal{Y}\rvert$
$1b$		challenge	$= 0$
$2a$	$D \in \neg\mathsf{SUC} \cap \neg\mathsf{CL} \cap \mathsf{SZ}_{\leq s}$	commitment	$\leq s \cdot \ell/\lvert\mathcal{Y}\rvert$
$2b$		challenge	$\leq \gamma_{cl}(sz^{\mathfrak{S}}_{triv})^N/\lvert\mathcal{Y}\rvert$

convention of $D(x) = \perp$ implies that $D \in \mathsf{CL} \Rightarrow D[x \mapsto U] \in \mathsf{CL}$ for all $U \in \mathcal{Y}$. Therefore, with the same argument as before, we may assume that $D \notin \mathsf{CL}$, as we could set $L^{x,D} = P|_{D^{|x}}$ which is empty.

1.) Let $D \in \mathsf{SUC} \cap \neg\mathsf{CL} \cap \mathsf{SZ}_{\leq s}$. We assume $D(x) = \perp$, which implies that $\perp \in \mathsf{P}'_{D^{|x}}$. Following Theorem 2, we must set $L^{x,D} := \mathsf{P}_{D^{|x}}$. As $D \in \mathsf{SUC}$, there are a $(\boldsymbol{y}_i)_{i\in[N']}$ and μ s.t.

$$\mathsf{V}(\mathsf{inst}_i, c_i, \boldsymbol{m}_{i,c_i}) \wedge (\mathsf{inst}, \mathsf{Ext}^*(\mathsf{inst}_i, \boldsymbol{m}_i)) \notin R$$

where $c_{i+1} := \gamma \circ D(i+1, \rho', \boldsymbol{y}_i, \mu) \wedge \boldsymbol{m}_i := D^{-1}(\boldsymbol{y}_i)$.

(a) Let x describe a commitment query to D with input m, i.e. $x = m$. By assumption, $D(m) = \perp$, and so, $D(m)$ is not a coordinate of any $\boldsymbol{y}_i \in \mathcal{Y}^\ell$ which we express here as $D(m) \notin \boldsymbol{y}_i$ for ease of notation. Furthermore,

$$U \in L^{x,D} \Leftrightarrow D[x \mapsto U] \in \mathsf{P} \Rightarrow D[x \mapsto U] \notin \mathsf{SUC} \Rightarrow \exists i \in [N'] : U \in \boldsymbol{y}_i,$$

where the last property can be seen by contraposition: Assuming that $U \notin \boldsymbol{y}_i$, then the defining properties of SUC remain as before. This means that

$$\Pr\left[U \in L^{x,D}\right] \leq \Pr[\exists i \in [N'] : U \in \boldsymbol{y}_i] \leq N \cdot \ell/\lvert\mathcal{Y}\rvert.$$

(b) Let x describe a challenge query. Because x describes a challenge query, it has no impact on preimages of $(\boldsymbol{y}_i)_{i\in[N]}$ and μ. They remain as they are, and hence $D[x \mapsto U] \in \mathsf{SUC}$ for any U which means that

$$\Pr\left[U \in L^{x,D}\right] = 0.$$

2.) Let $D \in \neg\mathsf{SUC} \cap \neg\mathsf{CL} \cap \mathsf{SZ}_{\leq s}$. We assume $D(x) = \perp$ which implies that $\perp \in \mathsf{P}_{D^{|x}}$. Following Theorem 2, we must set $L^{x,D} := \mathsf{P}'_{D^{|x}}$.

(a) Let x be a commitment query to D with input m, i.e. $x = m$. We then have

$$U \in L^{x,D} \Leftrightarrow D[x \mapsto U] \in \mathsf{P}' \Rightarrow \exists \boldsymbol{y} : D(\beta, \rho', \boldsymbol{y}, \mu) \neq \perp \wedge U \in \boldsymbol{y}$$

for some $\beta \in [N']$ and μ. The final implication can be seen by contradiction. Assume that $U \notin \boldsymbol{y}$, but $D[x \mapsto U] \in \mathsf{SUC}$. This implies that for every instance satisfying the SUC property, m can not be a preimage, because for a valid instance it must hold that each $c_{\beta+1} \neq \perp$. Therefore, the extraction would have also worked before setting x in the database, which means that $D \in \mathsf{SUC}$ already

forms a contradiction. Thus, for a transition, it must hold that there exists a $\boldsymbol{y}$ such that $D(\beta, \rho', \boldsymbol{y}, \mu) \neq \perp$ with $U \in \boldsymbol{y}$ in the database. By $\mathsf{SZ}_{\leq s}$ there are at most s queries with ℓ possible $y_i \in \mathcal{Y}$ as targets. Hence,

$$\Pr\left[U \in L^{x,D}\right] \leq \Pr[\exists \boldsymbol{y}, \mu, \beta : D(\beta, \rho', \boldsymbol{y}, \mu) \neq \perp \wedge U \in \boldsymbol{y}] \leq s \cdot \ell \cdot |\mathcal{Y}|^{-1}.$$

(b) Let x describe a challenge query, i.e. $x = (\beta, \rho', \boldsymbol{y}, \mu)$. We then have

$$\begin{aligned} &U \in L^{x,D} \Leftrightarrow D[x \mapsto U] \in P' = SUC \\ &\quad \Rightarrow \exists \boldsymbol{y}_i \in \mathcal{Y}^k \text{ for } i \in [N'] \setminus \{\beta - 1\} \text{ and } \boldsymbol{y}_{\beta-1} := \boldsymbol{y} \text{ s.t.} \\ &\quad \forall j \in [N'] : \mathsf{V}(\mathsf{inst}_j, c_j, \boldsymbol{m}_{j,c_j}) \wedge (\mathsf{inst}_j, \mathsf{Ext}(\mathsf{inst}_j, \boldsymbol{m}_j)) \notin R \end{aligned}$$

where $\boldsymbol{c}_j := D(j, \rho', \boldsymbol{y}_{j-1}, \mu)$ and $\boldsymbol{m}_j = D^{-1}(\boldsymbol{y}_j)$. The last implication can be seen by observing that because x is a challenge query, if it is not part of the proof satisfying the database property, the property would hold for D, too.

Considering the satisfying instance, we can observe that $\boldsymbol{y}$ uniquely defines $\boldsymbol{m}_{\beta-1}$. For these, we consider the set of possible challenges $S_{\beta-1} := \{c \in \mathcal{C} | \mathsf{V}(\mathsf{inst}_{\beta-1}, c, \boldsymbol{m}_{\beta-1,c})\}$. This set must be bounded by $sz^{\mathfrak{S}}_{triv}$, because there would be an extractor otherwise that could extract a witness from $m_{\beta-1}$ which would in turn mean, that m_β can not be contained in the instance satisfying SUC causing a contradiction. Each of these challenges c has at most one preimage, because there are no collisions in our database. This continues around the ring until we reach β. Here, there is again such a set S_β. Now we consider all possible paths and thereby get an upper bound on the number of possible values U can take on while causing the transition.

$$\Pr\left[U \in L^{x,D}\right] \leq \frac{\gamma_{cl} \cdot \prod_{i \in [N']} |S_i|}{|\mathcal{Y}|} \leq \frac{\gamma_{cl} \cdot (sz^{\mathfrak{S}}_{triv})^N}{|\mathcal{Y}|}.$$

Applying Theorem 2, we now get

$$[\![\mathsf{SZ}_{\leq s} \setminus \mathsf{SUC} \setminus \mathsf{CL} \to \mathsf{SUC}]\!] \leq \sqrt{10}\sqrt{|\mathcal{Y}|^{-1} \max\left\{N \cdot \ell, s \cdot \ell, \gamma_{cl} \cdot (sz^{\mathfrak{S}}_{triv})^N\right\}}.$$

□

C History-Freeness for Ring Signatures

We now formulate a version of history-freeness [12, Definition 4] for ring signatures. Additionally, we modify and generalize it in the following ways. First, we remove the pseudorandom function (PRF) used in the original version of [12, Definition 4] and use $2q$-wise independent functions [65] instead. Second, we also add explicit bounds to the final security statement.

Definition 22 (Problem [12, Definition 2]). *A problem P is defined by a pair $(\mathsf{Game}_P, \alpha_P)$ where Game_P specifies a game that a (possibly quantum) adversary plays with a classical challenger. The game works as follows:*

- *On input* 1^λ*, the challenger computes a value* x*, which it sends to the adversary as its input.*
- *The adversary is then run on* x*, and is allowed to make classical queries to the challenger.*
- *The adversary then outputs a value* y*, which it sends to the challenger.*
- *The challenger then looks at* x, y*, and the classical queries made by the adversary, and outputs* 1 *or* 0*.*

Here, $\alpha_P \in [0, 1)$*. It may also depend on* n*.*

We say that an adversary $\mathcal{A}$ wins the game Game_P if the challenger outputs 1. We define the advantage $\mathsf{Adv}_P^{\mathcal{A}}(\lambda)$ of $\mathcal{A}$ in problem P as

$$\mathsf{Adv}_P^{\mathcal{A}}(\lambda) = |\Pr[\mathcal{A}\,\text{wins in game}\,\mathsf{Game}_P] - \alpha_P|\,.$$

Definition 23 (Hard Problem [12, Definition 3]). *A problem* P *is hard for quantum computers, if for all polynomial time quantum adversaries* $\mathcal{A}$*,* $\mathsf{Adv}_P^{\mathcal{A}}(\lambda)$ *is negligible.*

The following definition is a direct adaptation of history-freeness for signatures [12, Definition 4]. The original definition only allowed for statistical distance as a measure between the actual sign algorithm and the simulated one. We observe that the Rényi divergence can also be used, and added this to our definition as an alternative.

Definition 24 (History-free Reduction for Ring Signatures). *A reduction from a hard problem* $P = (\mathsf{Game}_P, 0)$ *to the* UF-CRA1 *security of a ring signature* $\mathsf{RSig} = (\mathsf{Stp}, \mathsf{KGen}, \mathsf{Sign}^{\mathsf{O}}, \mathsf{Vf}^{\mathsf{O}})$ *in the ROM is called history-free, if it has the following form. Let* $\mathcal{A}$ *be an adversary against the* UF-CRA1 *security of* RSig *and* $\mathcal{B}$ *be the algorithm* $\mathcal{B}$ *for problem* P *such that:*

- *Algorithm* $\mathcal{B}$ *for* P *contains four (explicit, classical) algorithms:* START*,* $\mathsf{RAND}^{\mathsf{O}_c}$*,* $\mathsf{SIGN}^{\mathsf{O}_c}$*, and* $\mathsf{FINISH}^{\mathsf{O}_c}$*. The latter three algorithms have access to a shared random oracle* O_c*. These algorithms, except for* $\mathsf{RAND}^{\mathsf{O}_c}$*, may also make queries to the challenger for problem* P*. The algorithms are used as follows:*
  1. *Given an instance* x *for problem* P *as input, algorithm* B *first runs* $\mathsf{START}(x, N)$ *to obtain* (pk_i, z_i) *for* $i \in [N]$*,* pk_i *is a public key, and* z_i *is a private state.* $\mathcal{B}$ *gets these where* $z = (z_1, \ldots, z_N)$*, and sends* $\rho = \{\mathsf{pk}_1, \ldots, \mathsf{pk}_N\}$ *to* $\mathcal{A}$ *and plays the role of the challenger for* $\mathcal{A}$*.*
  2. *When* $\mathcal{A}$ *makes a random oracle query to* $\mathsf{O}(\rho', r)$*, algorithm* $\mathcal{B}$ *responds with* $\mathsf{RAND}^{\mathsf{O}_c}(\rho', r, z)$*. Note that* RAND *is given the current query as input, but is unaware of previous queries and responses.*
  3. *When* $\mathcal{A}$ *makes a signature query* $\mathsf{Sign}(\rho', \mathsf{sk}_i, m)$*, algorithm* $\mathcal{B}$ *responds with* $\mathsf{SIGN}^{\mathsf{O}_c}(\rho', m, z)$*.*
  4. *When* $\mathcal{A}$ *outputs a signature forgery* (m, σ) *for* $\rho^* \subseteq \rho$*, algorithm* $\mathcal{B}$ *outputs* $\mathsf{FINISH}^{\mathsf{O}_c}(\rho^*, m, \sigma, z)$*.*

- *There is an efficiently computable function* $\mathsf{INSTANCE}(\rho)$ *which produces an instance* x *of problem* P *such that* $\mathsf{START}(x) = (\rho, z)$ *for some* z*. Consider the process of first generating* $\mathsf{p} \leftarrow \mathsf{Stp}(1^\lambda)$*, then* $(\mathsf{sk}_i, \mathsf{pk}_i)$ *from* KGen *where* $\rho = \{(\mathsf{sk}_i, \mathsf{pk}_i)\}_{i\in[N]}$*, and then computing* $x = \mathsf{INSTANCE}(\rho)$*. The distribution of* x *generated in this way is* $\epsilon_{\mathsf{INSTANCE}}$ *close to the distribution of* x *generated in* Game_P*.*
- *For fixed* z*, consider the (classical) random oracle* $\mathsf{O}(\rho', r) = \mathsf{RAND}^{\mathsf{O}_c}(\rho', r, z)$*. Define a quantum oracle* O_{quant}*, which transforms a basis element* $|x, y\rangle$ *to* $|x, y \oplus \mathsf{O}(x)\rangle$*. The oracle, although noted slightly differently, is the same, remains the same. We require that* O_{quant} *is quantum computationally indistinguishable from a random oracle. The distinguishing advantage for* q_H *queries is bounded by* $\epsilon_{\mathsf{RAND},q_\mathsf{H}}$*.*
- $\mathsf{SIGN}^{\mathsf{O}_c}$ *either aborts (and hence* B *aborts) or it generates a valid signature relative to the oracle* $\mathsf{O}(\rho', r) = \mathsf{RAND}^{\mathsf{O}_c}(\rho', r, z)$ *with a distribution that for* q_s *sign queries is* either $\epsilon_{\mathsf{SIGN},q_s}$ *close to the correct signing algorithm or it has a small Rényi divergence* $\delta_{\mathsf{SIGN},q_s,\alpha}$ *compared to the correct signing algorithm.*[16] *The probability that none of the signature queries abort is at least* $1 - \epsilon_{abort}$*.*
- *If* (m, σ) *is a valid signature forgery relative to* ρ^* *and oracle* $\mathsf{O}(\rho^*, r) = \mathsf{RAND}^{\mathsf{O}_c}(\rho^*, r, z)$ *then the output of* $\mathcal{B}$ *(i.e.* $\mathsf{FINISH}^{\mathsf{O}_c(\rho^*,m,\sigma,z)}$*) causes the challenger for problem* P *to output* 1 *with probability at least* $1 - \epsilon_{\mathsf{FINISH}}$*.*

The following theorem is adapted from the corresponding theorem for signatures [12, Theorem 1]. We give a full proof and fill in the details needed for a concrete analysis of the security loss.

Theorem 12. *Let* $\mathsf{RSig} = (\mathsf{KGen}, \mathsf{Sign}, \mathsf{Vf})$ *be a ring signature scheme. Suppose that there is a history-free reduction for* RSig *with* $\epsilon_{\mathsf{INSTANCE}}$*,* $\epsilon_{\mathsf{RAND},q_\mathsf{H}}$*,* $\epsilon_{\mathsf{SIGN},q_s}$*,* $\delta_{\mathsf{SIGN},q_s,\alpha}$*,* ϵ_{abort}*,* $\epsilon_{\mathsf{FINISH}}$*. Then, for every* UF-CRA1 *adversary* $\mathcal{A}$ *against* RSig *making* q_s *sign queries and* q_H *queries to the QROM, we can construct a quantum adversary* $\mathcal{B}$ *for problem* P *such that*

$$\begin{aligned}\mathsf{Adv}^{\text{UF-CRA1}}_{\mathcal{A},\mathsf{RSig}}(\lambda) \leq \Pr[\mathcal{B}\, solves\, P] + \epsilon_{\mathsf{RAND},q_\mathsf{H}} + \epsilon_{\mathsf{INSTANCE}} \\ + \epsilon_{\mathsf{SIGN},q_s} + \epsilon_{abort} + \epsilon_{\mathsf{FINISH}}.\end{aligned}$$

and using the Rényi divergence, we can get

$$\begin{aligned}\mathsf{Adv}^{\text{UF-CRA1}}_{\mathcal{A},\mathsf{RSig}}(\lambda) \leq \left[\delta_{\mathsf{SIGN},q_s,\alpha}\left(\Pr[\mathcal{B}\, solves\, P] + \epsilon_{\mathsf{INSTANCE}} + \epsilon_{abort} + \epsilon_{\mathsf{FINISH}}\right)\right]^{\frac{\alpha-1}{\alpha}} \\ + \epsilon_{\mathsf{RAND},q_\mathsf{H}}\end{aligned}$$

Proof. The history-free reduction includes five (classical) algorithms START, RAND, SIGN, FINISH, and $\mathsf{INSTANCE}$. We consider a sequence of games where the first game is the standard UF-CRA1 ring signature game with respect to RSig in the QROM.

[16] The statement about the Rényi divergence was added to capture proofs that use a conceptually different proof technique requiring the Rényi divergence.

$\mathsf{Game}_1(\lambda)$**:** The quantum adversary $\mathcal{A}_Q$ plays the standard UF-CRA1 game for RSig. Assume that $\mathcal{A}_Q$ has a non-negligible winning probability.

$\mathsf{Game}_2(\lambda)$**:** Perform the following modifications compared to Game_1. After the challenger generates $(\mathsf{pk}_i, \mathsf{sk}_i) \leftarrow \mathsf{KGen}(\mathsf{p})$ for $i \in [N]$ where $\mathsf{p} \leftarrow \mathsf{Stp}(1^\lambda)$, it computes $x \leftarrow \mathsf{INSTANCE}(\rho)$ as well as $(\rho, z) \leftarrow \mathsf{START}(x)$. Further, instead of answering $\mathcal{A}_Q$'s quantum random oracle queries with a truly random oracle, the challenger simulates for $\mathcal{A}_Q$ a quantum accessible oracle O_q as an oracle that maps a basis element $|x, y\rangle$ into the element $|x, y \oplus \mathsf{RAND}^{\mathsf{O}_q(\rho', x, z)}\rangle$, where O_q is a random quantum-accessible oracle. By Definition 24, O_{quant} is computationally indistinguishable from random for quantum adversaries. The success probability of $\mathcal{A}_Q$ is thus $\epsilon_{\mathsf{RAND},q_\mathsf{H}}$ close to the probability in Game_1,

$$|\Pr[\mathcal{A}_Q \text{ wins in } \mathsf{Game}_1] - \Pr[\mathcal{A}_Q \text{ wins in } \mathsf{Game}_2]| \leq \epsilon_{\mathsf{RAND},q_\mathsf{H}}.$$

$\mathsf{Game}_3(\lambda)$**:** Modify the challenger from Game_2 as follows: instead of generating $(\mathsf{sk}_i, \mathsf{pk}_i)$ and computing $x = \mathsf{INSTANCE}(\rho)$, start off by running the challenger for problem P. When the challenger sends x, then start the challenger from the previous game using this x. The challenger is no longer in possession of sk_i, therefore, when $\mathcal{A}_Q$ asks for a signature on m, answer with $\mathsf{SIGN}^{\mathsf{O}_q}(\rho', m, z)$. First, since $\mathsf{INSTANCE}$ is part of a history-free reduction, this change in how we compute x only affects the distribution of x with $\epsilon_{\mathsf{INSTANCE}}$, and hence the behavior of $\mathcal{A}_Q$. Second, as long as all signing algorithms succeed, changing how we answer signing queries only negligibly affects the behavior of $\mathcal{A}_Q$, i.e., either multiplicatively using the Rényi divergence $\delta_{\mathsf{SIGN},q_s,\alpha}$ and Lemma 1, or additively using the statistical distance $\epsilon_{\mathsf{SIGN},q_s}$.

$$\Pr[\mathcal{A}_Q \text{ wins in } \mathsf{Game}_2] \leq \epsilon_{\mathsf{SIGN},q_s} + \Pr[\mathcal{A}_Q \text{ wins in } \mathsf{Game}_3] + \epsilon_{\mathsf{INSTANCE}}$$

or alternatively using the Rényi divergence

$$\Pr[\mathcal{A}_Q \text{ wins in } \mathsf{Game}_2] \leq (\delta_{\mathsf{SIGN},q_s,\alpha}(\Pr[\mathcal{A}_Q \text{ wins in } \mathsf{Game}_3] + \epsilon_{\mathsf{INSTANCE}}))^{\frac{\alpha-1}{\alpha}}$$

For the Rényi divergence, it should be noted that the modification of the instance, i.e., removing the use of the secret key, is only possible after the signature queries are simulated. Otherwise, the queries cannot be answered.

$\mathsf{Game}_4(\lambda)$**:** Define Game_4 as in Game_3, except for the following: First, generate a key k for the $2(q_\mathsf{H}+q_s+1)$-wise independent function F[17], and then to answer a query to O_q, the challenger applies the unitary transformation that takes a basis element $|x, y\rangle$ into $|x, y \oplus F(\rho', k, x)\rangle$. This can perfectly simulate the random oracle [65, Theorem 6.1], i.e.

$$\Pr[\mathcal{A}_Q \text{ wins in } \mathsf{Game}_3] = \Pr[\mathcal{A}_Q \text{ wins in } \mathsf{Game}_4]$$

[17] [12] use quantum-accessible PRFs, but Zhandry [65, Theorem 6.1] showed that $2(q_\mathsf{H}+q_s+1)$-wise independent functions suffice in this step.

Given a quantum adversary that has a non-negligible advantage in Game_4, we construct a quantum algorithm $\mathcal{B}_Q$ that solves the problem P. When $\mathcal{B}_Q$ receives instance x from the challenger for problem P, it computes $(\rho, z) \leftarrow \mathsf{START}(x, N)$ and generates a key k for F. Then, it simulates $\mathcal{A}_Q$ on ρ. $\mathcal{B}_Q$ answers $\mathcal{A}_Q$'s quantum random oracle queries using $\mathsf{RAND}^{F(\rho',k,\cdot)}(\rho', \cdot, z)$. It answers signing queries using $\mathsf{SIGN}^{F(\rho',k,\cdot)}(\rho', \cdot, z)$. Then, when $\mathcal{A}_Q$ outputs a forgery candidate (m, σ) for ρ^*, $\mathcal{B}_Q$ computes $\mathsf{FINISH}^{F(\rho^*,k,\cdot)}(\rho^*, m, \sigma, z)$, and returns the result to the challenger for problem P.
Observe that the behavior of $\mathcal{A}_Q$ in Game_4 is identical to that of $\mathcal{A}_Q$ as a subroutine of $\mathcal{B}_Q$. If (m, σ) is a valid forgery, then since FINISH is part of a history-free reduction, $\mathsf{FINISH}^{F(\rho^*,k,\cdot)}(\rho^*, m, \sigma, z)$ will cause the challenger for the problem P to accept with non-negligble probability contradicting our assumption that P is hard for quantum computers. Now, with probability $(1 - \epsilon_{abort})$, no abort occurs, and with probability at least $(1 - \epsilon_{\mathsf{FINISH}})$, the output is a valid solution for the problem P. Combined, this yields an adversary $\mathcal{B}_Q$

$$\Pr[\mathcal{A}_Q \text{ wins in } \mathsf{Game}_4] \leq \Pr[\mathcal{B}_Q \text{ solves } P] + \epsilon_{abort} + \epsilon_{\mathsf{FINISH}}.$$

□

C.1 Proof of Theorem 7

Proof (Theorem 7). By adapting the proof of [12, Theorem 2] we can construct a history-free reduction for RSig. Using Theorem 12, a generalization of the QROM lifting theorem for history-free reductions [12, Theorem 1], we then obtain a QROM security proof. We analyze the explicit security bounds this chain of arguments supports. We now describe the history-free reduction to the collision-resistance of the RPSF, i.e., the FINISH algorithm outputs a probable collision.

- On input ρ, define $\mathsf{INSTANCE}(\rho) := \rho$ and $\mathsf{START}(\rho) = (\rho, \rho)$. The key generation of the ring signature is exactly the key generation of the RPSF which refers to the hard problem of collision-resistance of the RPSF we reduce to.
- When $\mathcal{A}$ queries the (hash) oracle $\mathsf{H}(\rho', r)$, $\mathcal{A}_{\mathrm{COL}}$ responds with

$$\mathsf{RAND}^{\mathsf{O}_c}(\rho', r, \rho) := f_{\rho'}(\mathsf{SampleDom}_{\rho'}(1^\lambda : \mathsf{O}_c(\rho', r))).$$

- When $\mathcal{A}$ queries for a signature on (ρ', m) signed by party i where $\mathsf{pk}_i \in \rho' \cap \rho$, $\mathcal{A}_{\mathrm{COL}}$ responds with

$$\mathsf{SIGN}^{\mathsf{O}_c}(\rho', m, \rho) := (\mathsf{SampleDom}_{\rho'}(1^\lambda : \mathsf{O}_c(\rho', m))).$$

- When $\mathcal{A}$ outputs (ρ^*, m^*, σ^*), $\mathcal{A}_{\mathrm{COL}}$ outputs a potential collision

$$\mathsf{FINISH}^{\mathsf{O}_c}(\rho^*, m^*, \sigma^*) := (\mathsf{SampleDom}_{\rho^*}(1^\lambda : \mathsf{O}_c(\rho^*, m^*)), \sigma^*).$$

The key generation of the ring keys remains the same as in the original game, which is the same as for the underlying RPSF, which means that the instances are distributed the same, i.e., $\epsilon_{\mathsf{INSTANCE}} = 0$.

By definition of a RPSF, the distribution of $f_{\rho'}(\cdot)$ is within TV distance $\epsilon^{dom}_{\mathsf{RPSF}}$ from uniform if at least one of the public keys in ρ' was generated honestly. As $\mathsf{O}(\rho', r) = \mathsf{RAND}^{\mathsf{O}_c}(\rho', r, \rho) = f_{\rho'}(\mathsf{SampleDom}_{\rho'}(1^\lambda : \mathsf{O}_c(\rho', r)))$ and O_c is a random oracle, the outputs $\mathsf{O}(\rho', r)$ are distributed independently according to a distribution that is at most $\epsilon^{dom}_{\mathsf{RPSF}}$ away from uniform. Using Theorem 3, for any algorithm $\mathcal{A}$ making q quantum queries to its oracle, the TV distance between the output of $\mathcal{A}$ using a random oracle or O is at most $8(q_s + q_{\mathsf{H}})\sqrt{\epsilon^{dom}_{\mathsf{RPSF}}} =: \epsilon_{\mathsf{RAND},q_{\mathsf{H}}}$.

SIGN is now following the distribution of conditional preimage sampling, which by definition of the RPSF is similar to normal signing up to a statistical distance of $q_s\epsilon^{pre}_{\mathsf{RPSF}} =: \epsilon_{\mathsf{SIGN},q_s}$ and a Rényi divergence of $(\delta^{\alpha\text{-}pre}_{\mathsf{RPSF}})^{q_s} =: \delta_{\mathsf{SIGN},q_s,\alpha}$ for q_s sign queries. In this process of replacement, there is an additional loss: the SIGN algorithm is now deterministic. This could be used for distinguishing. However, we are in the UF-CRA1 game, so the adversary can only query the signing oracle for a fixed ring and message once, so the determinism can not be detected.

Observe that RAND and SIGN correspond to each other in that for a given message m, they define the same value in the range, and implicitly in the domain. Furthermore, the algorithms abort in the same cases, e.g., if a signature is queried for a secret key sk that is not contained in the ring, or if the challenger does not have knowledge of it. If $\mathcal{A}$ outputs a valid forgery (m^*, σ^*) for $\rho^* \subseteq \rho$, $\mathsf{FINISH}^{\mathsf{O}_c}(\rho^*, m^*, \sigma^*)$ produces a collision with probability at least $1 - 2^{-\beta(\lambda)}$ due to the min-entropy of the RPSF, i.e., $\epsilon_{\mathsf{FINISH}} = 2^{-\beta(\lambda)}$. (m^*, salt^*) cannot be part of a sign query, so $\mathsf{O}_c(\rho^*, \mathsf{salt}^* \| m^*)$ is independent of the forgery. The signature produced in FINISH by running SIGN is thus using the conditional preimage sampling for a fixed target for which the min-entropy statement is defined. Finally, we apply Theorem 12 to get the claimed security bound. □

D AOS Ring Signatures from Merkle Tree Based C&O Protocols

Theorem 13. *Let Σ be a $\mathfrak{S}^*$-sound Merkle tree C&O Σ-protocol with challenge space $\mathcal{C} = \mathcal{C}_\lambda$, $\ell = \ell(1^\lambda)$ commitments, $\mathsf{H} : \{0,1\}^* \to \mathcal{Y}$ be a hash function modelled as a QROM and $\gamma : \mathcal{Y} \to \mathcal{C}_\lambda$ be a function with $\gamma_{cl} := \max_{c \in \mathcal{C}} |\{y \in \mathcal{Y} \mid \gamma(y) = c\}|$. Set $\omega = \omega(\lambda) := \max_{c \in \mathcal{C}_\lambda} |c|$. Furthermore, let $N \in \mathbb{N}$ be fixed with $N(2\ell - 1) \leq 2(q_{\mathsf{H}} - 1)$. For any* UF-NRA *adversary $\mathcal{A}$ making at most q_{H} quantum queries to H, there exists an efficient adversary $\mathcal{B}$ against witness-recovery of the*

Σ *with*

$$\begin{aligned}\mathsf{Adv}^{\text{UF-NRA}}_{\mathcal{A},\mathsf{AOS}(\Sigma),N}(\lambda) &\leq N \cdot \mathsf{Adv}^{\text{W-REC}}_{\mathcal{B},\Sigma}(\lambda) + 2 \cdot N(\omega(\lambda)\log \ell + 1) \cdot |\mathcal{Y}|^{-1} \\ &\quad + q_{\mathsf{H}}^2 10/|\mathcal{Y}| \left(\sqrt{(q_{\mathsf{H}}-1)\cdot\gamma_{cl}} + \sqrt{\max\left\{2(q_{\mathsf{H}}-1), \gamma_{cl}\cdot(sz^{\mathfrak{S}}_{triv})^N\right\}}\right)^2. \\ &\leq N \cdot \mathsf{Adv}^{\text{W-REC}}_{\mathcal{B},\Sigma}(\lambda) + 2 \cdot N(\omega(\lambda)\log \ell + 1) \cdot |\mathcal{Y}|^{-1} \\ &\quad + q_{\mathsf{H}}^2 (q_{\mathsf{H}}-1) 20\gamma_{cl}/|\mathcal{Y}| \left(1 + 2(sz^{\mathfrak{S}}_{triv})^{N/2} + (sz^{\mathfrak{S}}_{triv})^N\right).\end{aligned}$$

Proof (Outline). We omit the full proof here, instead we briefly describe how our previous proof would change.

First, we consider how Lemma 7 changes. We must redefine the success game to use a different procedure to recompute the openings. As in [28], this is basically just a backwards traversal of the Merkle tree. The collision and size definition stay the same. Therefore, we can still utilize Lemma 8. This leaves us to consider how to adapt Lemma 10. The cases $1b$ and $2b$ remain as they are, and it remains to consider the cases $1a$ and $2a$. In $1a$, the number of possible targets for a collision grows, and there are a total of $N \cdot (2\ell - 1)$ targets now, as $(2\ell - 1)$ limits the number of elements of the Merkle tree. In $2a$, the number each query now fixes at most 2 other commitments (challenge queries fix 1, and commitment queries fix up to 2). Hence, there are at most $2s$ targets to consider.[18] Combining them bounds the transition capacity as

$$\leq q_{\mathsf{H}}\sqrt{\frac{10}{|\mathcal{Y}|}}\left(\sqrt{(q_{\mathsf{H}}-1)\cdot\gamma_{cl}} + \sqrt{\max\left\{N(2\ell-1), 2(q_{\mathsf{H}}-1), \gamma_{cl}\cdot(sz^{\mathfrak{S}}_{triv})^N\right\}}\right).$$

Then, we would use this transition bound in Theorem 10. Due to the switch, we have to make the same modifications that were done in [28]. More precisely, this means that we can keep almost everything as it is, except that we must check for consistency of the eventual signatures. This signature now builds a Merkle tree in its verification, and the verification of a response from the Σ-protocol requires $\log \ell$ many consistency checks (the height of the Merkle tree). Combining this with the ring size of N and the consistency between the classical queries in the verification, and the database is given except with probability

$$2 \cdot N(\omega(\lambda) \cdot \log \ell + 1) \cdot |\mathcal{Y}|^{-1}.$$

The remainder of the proof is analogous to the proof of Theorem 10. □

E Applications and Instantiations

This section serves to demonstrate how our techniques can be used in practice. We provide some background in the additional preliminaries in Appendix A.6

[18] In [28], the bound was given as $s \cdot (2\ell - 1)$, but it can be reduced to this tighter bound.

E.1 Falcon in the QROM

Recently, the security of Falcon has been revisited formally in [34] at Eurocrypt 2026, but their proof does not look at the QROM. [48] uses the Rényi divergence for a QROM proof of Falcon, but only to derive a bound on the statistical distance using Pinsker's inequality, which, as shown in [34], is already insufficient for the classical proof. To the best of our knowledge, there is no QROM proof for Falcon that avoids costly statistical distance arguments for the preimage sampler.

Falcon as a signature scheme can be classified as a Hash-then-Sign with aborts (HSwA) signature scheme. Aborts have been considered in the QROM both for signatures and sigma protocols, e.g., in [31,25,5]. The work of Fallahpour et al. [31] could be used for Falcon. We choose to write a direct proof, to avoid a separate analysis of the necessary accepting-transcript HVZK property using Rényi divergence arguments. The work of [25] is written for Fiat-Shamir with aborts (FSwA), but when considering generalised FSwA by fixing the first message of the sigma protocol as the randomness used in the signature, this can capture HSwA signatures. However, we were not able to directly map their simulator to Falcon.

The previous general solutions also share another problem that hinders their application: they assume that the challenge/hash value is distributed uniformly at random, which is not the case for Falcon. In Falcon, the values that are reprogrammed into the random oracle come from a distribution that is close to uniform in terms of Rényi divergence.

As no existing proof technique can be used directly, we translate the classical proof from [34] into the QROM.

To prove the security of Falcon in the QROM, we use the generalized abstraction of Falcon to CoreFalcon from [34] described in Fig. 9.

The proof we show here is primarily to demonstrate that Falcon can be proven secure in the QROM using Rényi divergence arguments that mirror the classical proof relatively closely. Compared to the detailed analysis done in [34], we have primarily converted the arguments to the QROM. One place where our proof departs from the ROM reasoning is when reducing to $\mathcal{R}$-NTRU-ISIS. In [34], a reduction to $\mathrm{q_H}+1$-$\mathcal{R}$-NTRU-ISIS is performed by programming a new instance for each query to the random oracle. This clearly does not work in the QROM. We provide two alternative ways out. i) We use measure-and-reprogram to include a single $\mathcal{R}$-NTRU-ISIS target, instead of their variant with $(q_{\mathsf{H}}+1)$ preimages. ii) We embed a multi-challenge t-$\mathcal{R}$-NTRU-ISIS challenge using small-range distributions from Lemma 16, yielding a reduction to t-$\mathcal{R}$-NTRU-ISIS for $t = \Omega\left(\frac{(q_{\mathsf{H}}+C_s)^3}{\epsilon}\right)$ with a success probability loss of ϵ and a longer reduction runtime.

Theorem 5 (SUF-CMA **of CoreFalcon$^+$**). *For any adversary $\mathcal{A}$ against the* SUF-CMA *security of CoreFalcon$^+$, making at most q_s signing queries and q_{H} QROM queries, there exists an adversary $\mathcal{B}$ against* $\mathrm{C_s}$-$\mathcal{R}$-NTRU-SPISIS*, *an adversary $\mathcal{C}$ against* $\mathcal{R}$-NTRU-ISIS *and an adversary $\mathcal{D}$ against* t-$\mathcal{R}$-NTRU-ISIS,

$\mathsf{Gen}(1^\lambda)$	$\mathsf{Vf}(\mathsf{pk}=\boldsymbol{h},\mu,\sigma=(r,\boldsymbol{s}_2))$
$(\boldsymbol{h},\boldsymbol{f},\boldsymbol{g}) \leftarrow \mathsf{TpdGen}_{\mathsf{NTRU}}(q,\alpha)$	$c := \mathsf{H}(\mathsf{pk}, r, \mu)$
$\boldsymbol{B} := \boldsymbol{B}_{\boldsymbol{f},\boldsymbol{g}}$	$\boldsymbol{s}_1 := \boldsymbol{c} - \boldsymbol{s}_2\boldsymbol{h} \mod q$
return $(\mathsf{sk} := \boldsymbol{B}, \mathsf{pk} := \boldsymbol{h})$	**return** $[\![\|(\boldsymbol{s}_1,\boldsymbol{s}_2)\|_2 \leq \beta]\!]$
$\mathsf{Sign}(\mathsf{sk}=\boldsymbol{B},\mu)$	$\mathsf{Sign}^+(\mathsf{sk}=\boldsymbol{B},\mu)$
$r \leftarrow \{0,1\}^k$	**while** $\|(\boldsymbol{s}_1,\boldsymbol{s}_2)\|_2 > \beta$:
$\boldsymbol{c} := \mathsf{H}(\mathsf{pk}, r, \mu) \in \mathcal{R}_q$	$r \leftarrow \{0,1\}^k$
while $\|(\boldsymbol{s}_1,\boldsymbol{s}_2)\|_2 > \beta$:	$\boldsymbol{c} := \mathsf{H}(\mathsf{pk}, r, \mu) \in \mathcal{R}_q$
$(\boldsymbol{s}_1,\boldsymbol{s}_2) \leftarrow \mathsf{PreSmp}_{\mathsf{NTRU}}(\boldsymbol{B}, s, \boldsymbol{c})$	$(\boldsymbol{s}_1,\boldsymbol{s}_2) \leftarrow \mathsf{PreSmp}_{\mathsf{NTRU}}(\boldsymbol{B}, s, \boldsymbol{c})$
$\sigma := (r,\boldsymbol{s}_2) \in \{0,1\}^k \times \mathcal{R}$	$\sigma := (r,\boldsymbol{s}_2) \in \{0,1\}^k \times \mathcal{R}$
return σ	**return** σ

Fig. 9. Variations of FALCON, namely CORFALCON = (Gen, Sign, Vf) and CORFALCON$^+$ = (Gen, Sign$^+$, Vf) from [34]. We used the notation of [35] for NTRU lattices, which differs slightly from [34], but they refer to the same underlying construction

s.t. for all $C_s, t \in \mathbb{N}^{\geq 1}, \alpha_u, \alpha_p \in \mathbb{R}^{>1}$*:*

$$\mathsf{Adv}^{\text{SUF-CMA}}_{\mathcal{A},\text{CoreFalcon}^+,q_s}(\lambda) \leq \left(r_u^{C_s}\left(r_p^{C_s}\left(\epsilon+\epsilon'\right)\right)^{\frac{\alpha_p-1}{\alpha_p}}\right)^{\frac{\alpha_u-1}{\alpha_u}} + \delta$$

where ϵ *and the runtimes of* $\mathcal{B}$ *and* $\mathcal{D}$ *are bounded in Table 1 and the other parameters, including* δ *as a function of* C_s *which is in a trade-off relationship with* t *via epsilon, are defined in Table 2.*

Proof. The proof follows the structure of [34, Theorem 1 and 2], and we only highlight those parts that are changed when transitioning to the QROM.

$\mathsf{Game}_1(\lambda)$**:** This is the unforgeability game for CORFALCON$^+$ so by definition

$$\Pr[1 \leftarrow \mathsf{Game}_1] = \mathsf{Adv}^{\text{SUF-CMA}}_{\mathcal{A},\text{CoreFalcon}^+,q_s}(\lambda).$$

We assume from here on, for the cost of one additional hash query, that the final forgery was queried for its random oracle value.

$\mathsf{Game}_2(\lambda)$**:** 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 the threshold C_s. Following [34, Claim 1.], we have

$$|\Pr[1 \leftarrow \mathsf{Game}_1] - \Pr[1 \leftarrow \mathsf{Game}_2]| \leq \sum_{i=0}^{q_s} \binom{C_s}{i} (1-p)^{C_s-i}(p)^i.$$

The signing oracle still runs $\mathsf{PreSmp}_{\mathsf{NTRU}}$ here, so p is the acceptance probability of the game we are leaving. The threshold is never reverted; all later games inherit the abort.

$\mathsf{Game}_3(\lambda)$**:** This game is identical to Game_2 except that we change how signature queries are answered. Instead of querying the random oracle, we program it with $\boldsymbol{c}$ from the freshly generated $\mathrm{C_s}$-$\mathcal{R}$-NTRU-SPISIS* samples, and we do so for every preimage sampling attempt, including those that the rejection loop later discards. Note that we think of this reprogramming as an operation applied to the random oracle, so when we later replace the random oracle by a different function, the reprogramming operation is applied to that oracle instead. The analysis of this step deviates from [34], as we need to account for the QROM. The total number of values whose distribution is changed is C_s, and there are a total of $C_s + q_{\mathsf{H}} + 1$ queries to the random oracle. As the challenge carries no rejection loop, each $\boldsymbol{c}_i$ is distributed exactly as $\mathcal{Q}_{\boldsymbol{h}}$. This new distribution replaces uniform values, and there are C_s independent replacements, which allows us to directly apply the maximal Rényi divergence over the choice of $(\boldsymbol{h}, \boldsymbol{f}, \boldsymbol{g}) \leftarrow \mathsf{TpdGen}_{\mathsf{NTRU}}(q, \alpha)$. Applying Lemma 6 with the Rényi divergence yields

$$\Pr[1 \leftarrow \mathsf{Game}_2] \leq (r_u^{C_s} \Pr[1 \leftarrow \mathsf{Game}_3])^{\frac{\alpha_u - 1}{\alpha_u}} + \frac{3C_s}{2}\sqrt{\frac{(C_s + q_{\mathsf{H}} + 1)}{2^k}}.$$

$\mathsf{Game}_4(\lambda)$**:** Conditioned on $\boldsymbol{c}_i$, the pair $(\boldsymbol{u}_i, \boldsymbol{v}_i)$ from the $\mathrm{C_s}$-$\mathcal{R}$-NTRU-SPISIS* challenge is distributed exactly as $\mathcal{D}_{\boldsymbol{\Lambda}_{\boldsymbol{c}_i}, s}$, since $(\boldsymbol{0}, \boldsymbol{c}_i)$ is a preimage of $\boldsymbol{c}_i$ and $\mathcal{D}_{\mathcal{R},s} \times \mathcal{D}_{\mathcal{R},s}$ restricted to that coset is the coset discrete Gaussian. This is precisely the distribution against which r_p is measured in Table 2. This game is identical to Game_3 except that the output of the preimage sampler is replaced by the pair $(\boldsymbol{s}_1, \boldsymbol{s}_2) := (\boldsymbol{v}_i, \boldsymbol{u}_i)$ generated in the previous game for computing $\boldsymbol{c}_i$, and the norm check is applied to that pair. This change in distribution behaves as in [34], giving us

$$\Pr[1 \leftarrow \mathsf{Game}_3] \leq (r_p^{C_s} \Pr[1 \leftarrow \mathsf{Game}_4])^{\frac{\alpha_p - 1}{\alpha_p}}.$$

Attempt i is now discarded exactly when $\|(\boldsymbol{u}_i, \boldsymbol{v}_i)\|_2 > \beta$, so the loop of Sign^+ is simulated faithfully, and Game_4 no longer uses the trapdoor. Let $(\mu^*, r^*, \boldsymbol{u}^*, \boldsymbol{v}^*)$ be the final forgery of an adversary, where $(\boldsymbol{u}^*, \boldsymbol{v}^*) := (\boldsymbol{s}_2^*, \boldsymbol{c}^* - \boldsymbol{s}_2^* \boldsymbol{h})$ for $\boldsymbol{c}^* := \mathsf{H}(\boldsymbol{h}, r^*, \mu^*)$. If (μ^*, r^*) was part of a programmed attempt i^*, then we have a solution for $\mathrm{C_s}$-$\mathcal{R}$-NTRU-SPISIS*, as the forgery is a different preimage of $\boldsymbol{c}_{i^*}$: if the attempt was accepted this follows from SUF-CMA freshness, and if it was discarded it follows from $\|(\boldsymbol{u}_{i^*}, \boldsymbol{v}_{i^*})\|_2 > \beta \geq \|(\boldsymbol{u}^*, \boldsymbol{v}^*)\|_2$.[19] Else, if (μ^*, r^*) has the same $\boldsymbol{c}$ as one of the programmed attempts, then $\mathcal{A}$ has hit one of the targets $T := \{\boldsymbol{c}_i\}_{i \in [C_s]}$ with an unreprogrammed hash value. The targets are drawn as $\mathrm{C_s}$-$\mathcal{R}$-NTRU-SPISIS*

[19] Note that the discarded case is only free because the challenge preimages are not conditioned on being short; otherwise it would require a conditional min-entropy argument.

challenges and are thus fixed independently of the oracle, and $|T| \leq C_s$, so their distribution is immaterial here. Substituting $\Pr[U \in T] \leq C_s/|\mathcal{R}_q|$ for $\Pr\left[U \in L^{x,D}\right]$ in the proof of Lemma 8 gives a transition capacity of $\sqrt{10C_s/|\mathcal{R}_q|}$ that is independent of the database size and holds at every step; summing over the $C_s + q_{\mathsf{H}} + 1$ queries and squaring as in Lemma 9 bounds this case by $\frac{10(C_s+q_{\mathsf{H}}+1)^2 C_s}{|\mathcal{R}_q|}$. Finally, there are two options to bound the last case, yielding the two different statements in our theorem.
The first approach uses measure-and-reprogram (Theorem 1) on the initial oracle before any other reprogramming. So, there are two levels of reprogramming, measure-and-reprogramming is the *inner* reprogramming and resampling is the *outer* reprogramming. In the case we are considering here, verifying the output forgery does not use a hash input with outer reprogramming, so the measure-and-reprogram experiment fails unless inner reprogramming happens for an input that is not subject to outer reprogramming. The freshly sampled output of measure-and-reprogram is now replaced by the challenge from a $\mathcal{R}$-NTRU-ISIS adversary. This captures all possible cases, which means

$$\begin{aligned}\Pr[1 \leftarrow \mathsf{Game}_4] \leq\ & \mathsf{Adv}^{\mathrm{C_s}\text{-}\mathcal{R}\text{-NTRU-SPISIS}^*}_{\mathcal{B},q,\alpha,\beta}(\lambda) \\ & + (2q_{\mathsf{H}} + 2C_s + 3)^2 \mathsf{Adv}^{\mathcal{R}\text{-NTRU-ISIS}}_{\mathcal{C},1,q,\alpha,\beta}(\lambda) \\ & + \frac{10(C_s + q_{\mathsf{H}} + 1)^2 C_s}{|\mathcal{R}_q|}.\end{aligned}$$

The second approach uses small-range distributions (Lemma 16). Given t samples from a t-$\mathcal{R}$-NTRU-ISIS adversary, we can replace the random oracle by a small-range function with range size t. Note that we replace the un-reprogrammed random oracle by a small-range distribution, effectively replacing the reprogrammed oracle with a small-range distribution that has been reprogrammed in the same way. This transition is possible as the un-reprogrammed oracle is indistinguishable from the small-range version, and the other game hops that happened after can be simulated by a distinguisher. Now, every valid forgery is a valid answer in the t-$\mathcal{R}$-NTRU-ISIS game, as a valid forgery will always relate to an un-reprogrammed hash output. This transition does incur an additive loss in the order of $\frac{27(q_{\mathsf{H}}+C_s+1)^3}{t}$ such that

$$\begin{aligned}\Pr[1 \leftarrow \mathsf{Game}_4] \leq\ & \mathsf{Adv}^{\mathrm{C_s}\text{-}\mathcal{R}\text{-NTRU-SPISIS}^*}_{\mathcal{B},q,\alpha,\beta}(\lambda) \\ & + \mathsf{Adv}^{\mathrm{t}\text{-}\mathcal{R}\text{-NTRU-ISIS}}_{\mathcal{D},1,q,\alpha,\beta}(\lambda) + \frac{27(q_{\mathsf{H}} + C_s + 1)^3}{t} \\ & + \frac{10(C_s + q_{\mathsf{H}} + 1)^2 C_s}{|\mathcal{R}_q|}.\end{aligned}$$

As the reduction needs to write the t $\mathcal{R}$-NTRU-ISIS samples into its QRAM to instantiate the small-range distribution, it has an additional additive runtime overhead of $\Theta(t)$ compared to the the measure-and-reprogram reduction.

□

Remark 3. The C_s-$\mathcal{R}$-NTRU-SPISIS* term is only needed for SUF-CMA. Every programmed position refers to a message that was submitted to the signing oracle, so under EUF-CMA the freshness condition on μ^* excludes that case and the corresponding summand can be dropped from ϵ'.

E.2 Application to GANDALF

This section performs a QROM security analysis of GANDALF using our techniques. The original ring signature scheme [35, Figure 5] can be translated into a RPSF (as detailed in Fig. 10).

When instantiating GANDALF as a RPSF-based ring signature (see Fig. 1), the resulting signature is syntactically equivalent to GANDALF, except for two small modifications. First, we need to add k salt bits such that we can use the Rényi divergence properties (see Section 5.1). This does not have any functional impact on the RPSF and is also what is done for GANDALF, when turning it into a SUF-CRA ring signature [35, Appendix B.2, Theorem 7]. Secondly, the RPSF presampling algorithm returns $N+1$ elements for a ring ρ with $|\rho| = N$ instead of N as in GANDALF. This one element can, however, be saved by changing the evaluation in Fig. 1 such that the RPSF recomputes the missing element with the hash target and then all are evaluated. Thus, this second modification is just an additional optimization in GANDALF and has no loss in security.

Properties of the RPSF The effective RPSF built from GANDALF has slightly different correctness properties than those outlined in the definition of RPSFs (Definition 1). Both $\mathsf{SampleDom}$ and $\mathsf{SamplePre}$ for a fixed ρ are defined over $\mathcal{R}^{|\rho|}$ instead of D_ρ which is a ball in $\mathcal{R}^{|\rho|}$. $\mathsf{SamplePre}$ will always return valid preimages under evaluation of the map, but the images might not be in the domain. The correctness property of RPSFs, for simplicity, does not consider that $\mathsf{SampleDom}$ and $\mathsf{SamplePre}$ might be outside the domain. To account for this, the actual function f_ρ needs to check if the provided value is in the domain when used in Fig. 1. In the RSig proof (Theorem 6), treat f_ρ simply as the map from $\mathcal{R}^{|\rho|}$ to $\mathcal{R}_q$. However, when the adversary provides a forgery for a previously signed target, this is only a proper signature, if the simulated signature is also in the domain. Therefore, we get an additional error term based on whether the sample from $\mathsf{SampleDom}$ is in the domain. This is bounded using Lemma 12 and describes the additional error term (multiplied by q_s).

Correctness The correctness can be imported directly from [35], as the two ring signatures are semantically equivalent.

Corollary 4. RPSF *(Fig. 10) is* $\delta(\kappa)$*-correct where*

$$\delta(\kappa) = \tau^{(\kappa+1)M} \cdot e^{\frac{(\kappa+1)M}{2}(1-\tau^2)}$$

with $\tau > 1$.

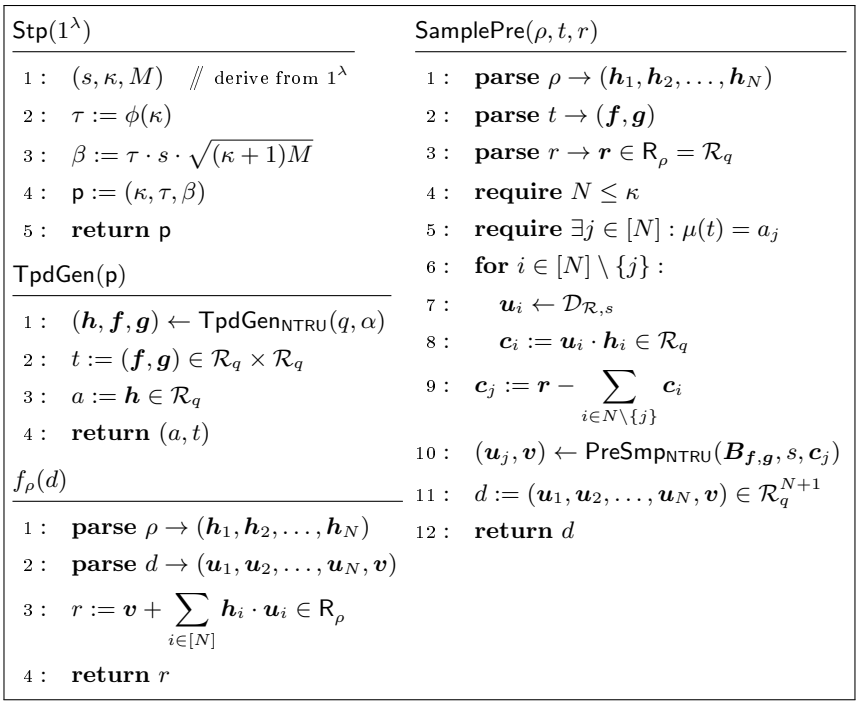

Fig. 10. A RPSF build from GANDALF. The domain D_ρ for $|\rho| = N$ is $\{(\boldsymbol{d}_1, \boldsymbol{d}_2, \ldots \boldsymbol{d}_N, \boldsymbol{d}) \in \mathcal{R}_q^{N+1} : \|(\boldsymbol{d}_1, \boldsymbol{d}_2, \ldots, \boldsymbol{d}_N, \boldsymbol{d})\|_2 \leq \beta\}$. The range $\mathsf{R}_\rho = \mathcal{R}_q$ for every ρ. The algorithms $\mathsf{TpdGen}_{\mathsf{NTRU}}$ and $\mathsf{PreSmp}_{\mathsf{NTRU}}$ are from Lemma 13 and Corollary 3. The function μ identifies the public value a that belongs to the trapdoor t. $\mathsf{SampleDom}(\rho)$ for $|\rho| = N$ is $\mathcal{D}_{\mathcal{R},s}^{N+1}$.

Domain Sampling with Uniform Output. We can use the Rényi uniformity arguments for NTRU from Corollary 2.

Corollary 5. *Let q be prime, $\alpha \in (1,\infty), \epsilon \in (0,1/2), s \geq \eta_{\epsilon(\boldsymbol{\Lambda}_{\boldsymbol{h},q})}$. The* SampleDom *algorithm of Fig. 10 satisfies domain sampling with uniform outputs (see notation in Property 1) with*

$$\delta_{\mathsf{RPSF}}^{\alpha\text{-}dom} \lesssim 1 + \frac{2\alpha\epsilon^2}{(1-\epsilon)^2}.$$

Proof. Considering the definition of Property 1, there is at least one honestly generated (non-zero) public value $\boldsymbol{h}_j$ in the ring $\rho = (a_1,\ldots,a_N) = (\boldsymbol{h}_1,\ldots,\boldsymbol{h}_N)$ for $j \in [N]$. Let $d = (\boldsymbol{u}_1,\ldots,\boldsymbol{u}_{N+1})$ be a sample from SampleDom. Fixing all $i \in [N] \setminus j$, and the respective $\boldsymbol{u}_i$ from SampleDom, they describe a simple shift in $\mathcal{R}_q$:

$$f_\rho(d) = \underbrace{\boldsymbol{u}_j\boldsymbol{h}_j + \boldsymbol{u}_{N+1}}_{:=\boldsymbol{x}} + \sum_{i\in[N]\setminus\{j\}} \boldsymbol{u}_i\boldsymbol{h}_i.$$

Applying Corollary 2 to $\boldsymbol{x}$ gives the Rényi divergence bound, which is preserved by the fixed shift applied after. □

Preimage Sampling is not detectable. To show that the preimage sampling is not detectable, we can use Corollary 3.

Corollary 6. *Let M be a positive integer, $\alpha > 1$ and $\epsilon \in (0,1/4)$. Then, with parameter choice according to Corollary 3 ($s \geq \eta_{\epsilon(\mathbb{Z}^{2M})} \cdot \|\boldsymbol{B}\|_{GS}$), the* SamplePre *algorithm of Fig. 10 has Kullback-Leibler divergence and Rényi divergence bounded as*

$$\epsilon_{\mathsf{RPSF}}^{KL\text{-}pre} \lesssim 2\epsilon^2 \quad \textit{and} \quad \delta_{\mathsf{RPSF}}^{\alpha\text{-}pre} \lesssim 1 + 2\alpha\epsilon^2.$$

Proof. Considering the definition of Property 2, there is at least one honestly generated (non-zero) public value $\boldsymbol{h}_j$ in the ring $\rho = (a_1,\ldots,a_N) = (\boldsymbol{h}_1,\ldots,\boldsymbol{h}_N)$ for $j \in [N]$. Observe that Corollary 3 can be applied for an arbitrary target. For a preimage sample $d = (\boldsymbol{u}_1,\ldots\boldsymbol{u}_{N+1}) \leftarrow \mathsf{SamplePre}(\rho,t,\boldsymbol{r})$ where $j \in [N]$ is the honest index for which the trapdoor is given. Then, treat all other indices $[N+1]\setminus\{j,N+1\}$ as a target shift that is sampled before using the same distribution. Now, they have both a fixed arbitrary target with the same distribution so far. Applying Corollary 3 yields the respective bounds. □

One-Wayness.

Lemma 17. *For any adversary $\mathcal{A} = (\mathcal{A}_1,\mathcal{A}_2)$ against the one-wayness (Property 3) of Fig. 10, there exists an adversary $\mathcal{B}$ such that*

$$\mathsf{Adv}^{\mathrm{OW}}_{\mathcal{A},\mathsf{RPSF}}(\lambda) \leq \mathsf{Adv}^{\mathcal{R}\text{-NTRU-ISIS}}_{\mathcal{B},\kappa,q,\alpha,\beta}(\lambda).$$

Proof. Assume an adversary $\mathcal{A}$ against the one-wayness of Fig. 10. We can make a straightforward reduction to the $\mathcal{R}$-NTRU-ISIS problem.

Consider $\mathcal{R}$-NTRU-ISIS game with N values $\boldsymbol{h}_i$, β as defined in the setup algorithm of the RPSF and a random target $\boldsymbol{r} \leftarrow \mathcal{R}_q$. All $\boldsymbol{h}_i$ will now be given to $\mathcal{A}$ as the public values and $\boldsymbol{r}$. $\mathcal{A}_1$ outputs (state, ρ') where $\rho' \subseteq \rho$. We define L as the set of indices such that $\{\rho_i\}_{i\in L} = \rho'$. Now, we give the target $\boldsymbol{r}$ to $\mathcal{A}_2$ with state and we will receive a preimage d of r. The preimage $d = (\boldsymbol{u}_i)_{i\in L\cup\{N+1\}}$ must be in the domain to be properly evaluated by $f_{\rho'}$, so $\left\|(\boldsymbol{u}_i)_{i\in L\cup\{N+1\}}\right\|_2 \leq \beta$ and $f_{\rho'}(d) = \boldsymbol{r}$. This preimage can be extended to a $\mathcal{R}$-NTRU-ISIS solution by setting the remaining positions to $\mathbf{0}$, i.e., output $(\boldsymbol{u}_i)_{i\in[N+1]}$ where $\boldsymbol{u}_j = \mathbf{0}$ for $j \notin L \cup \{N+1\}$. The evaluation remains the same, and the norm bounds have not changed, making it a valid $\mathcal{R}$-NTRU-ISIS solution. □

Conditional Preimage Min-Entropy.

Lemma 18. *If $s \geq 2\eta_\epsilon(\boldsymbol{\Lambda}_{\boldsymbol{h},q})$ and $\epsilon \in (0, 1/3)$, then the conditional min-entropy described in Property 4 of* SampleDom *as defined in Fig. 10 is at least $2M - 1$.*

Proof. The domain is defined by the ring size N of ρ, where all are honestly generated. We consider a subring $\rho' \subseteq \rho$. To be meaningful ρ' is of size $N \geq 1$. SampleDom samples a discrete Gaussian over $\mathcal{R}^{N+1}$. Consider the first $N-1$ to be fixed. This allows us to use the conditional min-entropy of the normal NTRU PSFs, so we can use the existing literature. So, for any value of the first $N-1$ elements, we can already get high min-entropy.

We use the arguments from [61, Proof of Property 4 in Theorem 4.2]. For any fixed target $\boldsymbol{t} \in \mathcal{R}_q$, the conditional distribution of $\boldsymbol{z} = (\boldsymbol{z}_1, \boldsymbol{z}_2) \leftarrow \mathcal{D}^2_{\mathcal{R},s}$, given $\boldsymbol{h}\boldsymbol{z}_1 + \boldsymbol{z}_2 = \boldsymbol{t}$ is exactly $\boldsymbol{c} + \mathcal{D}_{\boldsymbol{\Lambda}_{\boldsymbol{h},q},s,-\boldsymbol{c}}$ where $\boldsymbol{c} = (\mathbf{1}, \boldsymbol{t} - \boldsymbol{h})$[20]. Now, apply Lemma 11 to get the conditional min-entropy of at least $2M - 1$. □

Collision-Resistance.

Lemma 19. *For any adversary $\mathcal{A}$ against the collision-resistance (Property 5) of Fig. 10, there exists an adversary $\mathcal{B}$ such that*

$$\mathsf{Adv}^{\mathrm{COL}}_{\mathcal{A},\mathsf{RPSF}}(\lambda) \leq \mathsf{Adv}^{\mathcal{R}\text{-NTRU-SIS}}_{\mathcal{B},\kappa,q,\alpha,2\beta}(\lambda).$$

Proof. The collision resistance, similarly to the one-wayness, can be converted directly from a subring ρ' to the highest ring ρ, by adding $\mathbf{0}$ values. Two different short values d_1 and d_2, mapping to the same target, means that $d_1 - d_2$ (coordinate-wise) maps to $\mathbf{0}$ due to the linearity. This is directly a solution to $\mathcal{R}$-SIS and has length at most 2β. □

QROM bounds on GANDALF The general proof for RPSF considers perfect correctness, but for GANDALF we need to consider that signatures and domain samples might be too long. The transitioning from presampling to the simulation using SampleDom does not use correctness properties and is independent of those. The security proof could only fail in the last step when reducing to $\mathcal{R}$-NTRU-SIS,

[20] $\boldsymbol{c}$ is a preimage of $\boldsymbol{t}$, i.e., $\boldsymbol{h}\mathbf{1} - \boldsymbol{h} + \boldsymbol{t} = \boldsymbol{t}$.

because the simulated values might be too large. Other positions do not fail, because they do not explicitly use the correctness properties, and the transitions themselves are independent of correctness. With a probability of at least $1 - q_s \cdot \delta(\kappa)$, all are correct, and the reduction works. Now we combine all previous results with Theorem 6 to get the following.

Theorem 14. *Let $\mathcal{R} = \mathbb{Z}[X]/(X^M + 1)$ where M is a power-of-two, q prime, $\alpha_1, \alpha_2 \in (1, \infty), \epsilon \in (0, 1/4), s \geq 2\eta_\epsilon(\boldsymbol{\Lambda}_{\boldsymbol{h},q})$, $\kappa \in \mathbb{N}$, $k \in \mathbb{N}$ and $\tau > 1$. For any adversary $\mathcal{A}$ against the* SUF-CRA *security of* RSig *(Fig. 1) instantiated with* RPSF *(Fig. 10) with* $\mathsf{PreSmp}_{\mathsf{NTRU}}$ *and* $\mathsf{TpdGen}_{\mathsf{NTRU}}$ *making at most q_s signature and q_{H} QROM queries to* H*, there exists adversaries $\mathcal{B}$ and $\mathcal{C}$ such that*

$$\mathsf{Adv}^{\text{SUF-CRA}}_{\mathcal{A},\mathsf{RSig},N,q_s}(\lambda) \leq \left(\left(\delta^{\alpha\text{-}dom}_{\mathsf{RPSF}}\right)^{q_s} \left((\delta^{\alpha\text{-}pre}_{\mathsf{RPSF}})^{q_s} \cdot \epsilon' \right)^{\frac{\alpha_2 - 1}{\alpha_2}} \right)^{\frac{\alpha_1 - 1}{\alpha_1}} + \frac{3q_s}{2} \sqrt{\frac{q_s + q_{\mathsf{H}} + 1}{2^k}}$$

where $\epsilon' = \mathsf{Adv}^{\mathcal{R}\text{-NTRU-SIS}}_{\mathcal{B},\kappa,q,\alpha,2\beta}(\lambda) + \frac{20(q_s+q_{\mathsf{H}}+1)^2 q_s}{q^M} + (2q_s + 2q_{\mathsf{H}} + 3)^2 \mathsf{Adv}^{\mathcal{R}\text{-NTRU-ISIS}}_{\mathcal{C},\kappa,q,\alpha,\beta}(\lambda) + q_s \cdot \delta(\kappa)$ *and* $\delta(\kappa) \leq \tau^{(\kappa+1)M} \cdot e^{\frac{(\kappa+1)M}{2}(1-\tau^2)}$.

When instantiating, explicit Rényi divergence analysis can be done using for example [35, Lemma 9] or optimized analysis of Rényi divergence as in [35].

Comparison to Classical Proof The QROM result is very close to what can be expected when transitioning the classical results to the QROM. Guessing the index for inclusion comes at the additional multiplicative cost of all queries, as we need to rely on Theorem 1. The term for hash collisions suffers the same factor as we transition to the QROM and is standard. Finally, the tight adaptive reprogramming Lemma 6 only introduces an additive term that can be controlled entirely by the saltbits k. The remaining terms remain untouched, showing that the classical proof can be transferred to the QROM without significant overhead.