

A Concrete Treatment of Fiat-Shamir Signatures in the Quantum Random-Oracle Model

Eike Kiltz ¹

Vadim Lyubashevsky ²

Christian Schaffner ³

February 20, 2018

¹ Ruhr Universität Bochum

eike.kiltz@rub.de

² IBM Research – Zurich

vad@zurich.ibm.com

³ QuSoft and ILLC, University of Amsterdam

c.schaffner@uva.nl

Abstract

The Fiat-Shamir transform is a technique for combining a hash function and an identification scheme to produce a digital signature scheme. The resulting scheme is known to be secure in the random oracle model (ROM), which does not, however, imply security in the scenario where the adversary also has quantum access to the oracle. The goal of this current paper is to create a generic framework for constructing tight reductions in the QROM from underlying hard problems to Fiat-Shamir signatures.

Our generic reduction is composed of two results whose proofs, we believe, are simple and natural. We first consider a security notion (UF-NMA) in which the adversary obtains the public key and attempts to create a valid signature without accessing a signing oracle. We give a tight reduction showing that deterministic signatures (i.e., ones in which the randomness is derived from the message and the secret key) that are UF-NMA secure are also secure under the standard chosen message attack (UF-CMA) security definition. Our second result is showing that if the identification scheme is “lossy”, as defined in (Abdalla et al. Eurocrypt 2012), then the security of the UF-NMA scheme is tightly based on the hardness of distinguishing regular and lossy public keys of the identification scheme. This latter distinguishing problem is normally exactly the definition of some presumably-hard mathematical problem. The combination of these components gives our main result.

As a concrete instantiation of our framework, we modify the recent lattice-based Dilithium digital signature scheme (Ducas et al., TCHES 2018) so that its underlying identification scheme admits lossy public keys. The original Dilithium scheme, which is proven secure in the classical ROM based on standard lattice assumptions, has 1.5KB public keys and 2.7KB signatures. The new scheme, which is tightly based on the hardness of the Module-LWE problem in the QROM using our generic reductions, has 7.7KB public keys and 5.7KB signatures for the same security level. Furthermore, due to our proof of equivalence between the UF-NMA and UF-CMA security notions of deterministic signature schemes, we can formulate a new non-interactive assumption under which the original Dilithium signature scheme is also tightly secure in the QROM.

1 Introduction

FIAT-SHAMIR SIGNATURES FROM IDENTIFICATION PROTOCOLS. A canonical identification scheme [AABN02] is a three-move authentication protocol ID of a specific form. The prover (holding the secret-key) sends a commitment W to the verifier. The verifier (holding the public-key) returns a random challenge c . The prover sends a response Z . Finally, using the verification algorithm, the verifier accepts if the transcript (W, c, Z) is correct. The Fiat-Shamir transformation [FS87, AABN02] combines a canonical identification scheme ID and a hash function H to obtain a digital signature scheme $FS = FS[ID, H]$. The signing algorithm first iteratively generates a transcript (W, c, Z) , where the challenge c is derived via



Figure 1: Known security results of Fiat-Shamir signatures $FS = FS[ID, H]$ in the ROM. Solid arrows denote tight reductions, dashed arrows non-tight reductions.

$c := H(W \parallel M)$. Signature $\sigma = (W, Z)$ is valid if the transcript $(W, c := H(W \parallel M), Z)$ makes the verification algorithm accept. Lyubashevsky [Lyu09] further generalized this to the “Fiat-Shamir with aborts” transformation to account for aborting provers.

SECURITY OF FIAT-SHAMIR SIGNATURES IN THE ROM. Security of $FS[ID, H]$ in the ROM can be proved in two steps. Firstly, if the underlying identification scheme has statistical Honest-Verifier Zero-Knowledge (HVZK), then UnForgeability against Chosen Message Attack (UF-CMA) and UnForgeability against No Message Attack (UF-NMA) are tightly equivalent (UF-NMA security means that the adversary is not allowed to make any signing queries). Secondly, the Forking Lemma [PS00, BN06] (based on a technique called “rewinding”) is used to prove UF-NMA security in the random-oracle model (ROM) [BR93] from computational Special Soundness (SS). The latter part of the security reduction is non-tight and the loss in tightness is known to be inherent (e.g., [PV05, KMP16]).

LOSSY IDENTIFICATION SCHEMES. With the goal of constructing signature schemes with a tight security reduction and generalizing a signature scheme by Katz and Wang [KW03], AFLT [AFLT12] introduced the new concept of lossy identification schemes and proved that Fiat-Shamir transformed signatures have a tight security reduction in the ROM. A lossy identification scheme comes with an additional lossy key generator that produces a lossy public key, computationally indistinguishable from a honestly generated public key. Further, relative to a lossy public key the identification scheme has *statistical soundness*, i.e., not even an unbounded adversary can successfully impersonate a prover. Figure 1 summarizes the known security results of Fiat-Shamir signatures in the ROM.

QUANTUM RANDOM-ORACLE MODEL. Recently, NIST announced a competition with the goal to standardize new asymmetric encryption and signature schemes [NIS13] with security against quantum adversaries, i.e., adversaries equipped with a quantum computer. There exists a number of (sometimes only implicitly defined) canonical identification schemes (e.g., [Ste94, MV03, Lyu08, KTX08, Lyu09, AFLT12, GLP12, Lyu12, DDL13, DLP14, BG14, Lyu16, ABB⁺17, DLL⁺17]) whose security relies on the hardness of certain problems over lattices and codes, which are generally believed to resist quantum adversaries. Quantum computers may execute all “offline primitives” such as the hash function on arbitrary superpositions, which motivated the introduction of the quantum (accessible) random-oracle model (QROM) [BDF⁺11]. That is, in the UF-CMA security experiment for signatures in the QROM, an adversary has quantum access to a perfect hash function H and classical access to the signing oracle. Aiding in the construction of UF-CMA secure signatures with provable (post-quantum) security in the QROM is the main motivation of this paper.

SECURITY OF FIAT-SHAMIR SIGNATURES IN THE QROM. A number of recent works considered the security of Fiat-Shamir transformed signatures in the QROM. [BDF⁺11] proved a general result showing that if a reduction in the classical ROM is *history-free*, then it can also be carried out in the QROM. History-free reductions basically determine random oracle answers independently of the history of previous queries. For reductions that are not history-free, adaptive re-programming of the quantum random oracle is required which is problematic in the QROM: with one single quantum query to all inputs in superposition, an adversary might learn a superposition of all possible random oracle values which essentially means the reduction has to provide plausible values for the whole random oracle at this point. Hence, adaptive reprogramming in the QROM is difficult (but not impossible e.g., [BBBV97, ES15, Unr15]).

Unfortunately, the known random-oracle proofs of Fiat-Shamir signatures [PS00, AFLT12, KMP16] are not history-free. Beyond the general problem of adaptive re-programming, the classical proof [PS00] uses rewinding and the Forking Lemma, a technique that we currently do not know how to extend to the

quantum setting. Even worse, Ambanis et al. [ARU14] proved that Fiat-Shamir signatures cannot be proven secure in a black-box way by just assuming computational special soundness and HVZK (these two conditions are, on the other hand, sufficient for a proof in the classical ROM.)

To circumvent the above negative result, Unruh [Unr15] proposed an alternative Fiat-Shamir transformation with provable QROM security but the resulting signatures are considerably less efficient as they require multiple executions of the underlying identification scheme.

Alkim et. al [ABB⁺17] gave a concrete tight security reduction for a signature scheme, TESLA, in the QROM. TESLA is a concrete lattice-based digital signature scheme implicitly derived via the Fiat-Shamir transformation. Their QROM proof from the learning with errors (LWE) assumption adaptively re-programs the quantum random oracle using a technique from [BBBV97] and seems tailored to their particular identification protocol. As described in [ABB⁺17], the intuition behind the QROM security proof for TESLA comes from the fact that the underlying identification scheme is lossy. They leave it as an open problem to prove Fiat-Shamir signatures generically secure from lossy identification schemes.

Unruh [Unr17a] could prove (among other things) that identification schemes with HVZK and statistical soundness yield UF-CMA secure Fiat-Shamir signatures in the QROM when additionally assuming a “dual-mode hard instance generator” for generating key pairs of the identification scheme. The latter dual mode hard instance generator is very similar to lossy identification schemes. Whereas the original publication [Unr17a] only contains asymptotic proofs, a recently updated version of the full version [Unr17b] also provides concrete security bounds. Below, in Section 1.2, we will compare them with our bounds.

1.1 Our Results

This work contains a simple and modular security analysis in the QROM of signatures $\text{FS}[\text{ID}, \text{H}]$ obtained via the Fiat-Shamir transform with aborts [Lyu09] from any lossy identification scheme ID. We also consider the security of a deterministic variant $\text{DFS}[\text{ID}, \text{H}, \text{PRF}]$ with better tightness. DFS derives the randomness for signing deterministically using a pseudo-random function PRF. Our main security statements are summarized in Figure 2. Most importantly, if ID is a lossy identification scheme and has HVZK, then $\text{DFS}[\text{ID}, \text{H}, \text{PRF}]$ is tightly UF-CMA secure and $\text{FS}[\text{ID}, \text{H}]$ is (non-tightly) UF-CMA secure in the QROM. Our results suggest to prefer $\text{DFS}[\text{ID}, \text{H}, \text{PRF}]$ over $\text{FS}[\text{ID}, \text{H}]$.

The main component of our proof is a tweak to the AFLT Fiat-Shamir proof [AFLT12] that makes it history-free. Together with the general result of [BDF⁺11], one can immediately obtain *asymptotic* (i.e., non-concrete) versions of our QROM proof as a simple corollary. In this work, we instead give direct proofs with concrete, tight security bounds.

To demonstrate the efficacy of our generic framework, we construct a lattice-based signature scheme. The most compact lattice-based schemes, in terms of public key and signature sizes, crucially require sampling from a discrete Gaussian distribution [DDLL13, DLP14]. Such schemes, however, have been shown to be particularly vulnerable to side-channel attacks (c.f. [BHL16, EFGT17]), and it therefore seems prudent to consider schemes that only require simple uniform sampling over the integers. Of those, the most currently efficient one is the Dilithium signature scheme [DLL⁺17]. This signature scheme is proved secure based on the MSIS (Module-SIS) and the MLWE (Module-LWE) assumptions in the ROM implicitly using the framework from Figure 1.

In this paper, we provide a practical instantiation of a lossy identification scheme to obtain a new digital signature scheme, Dilithium-QROM, with a tight security reduction in the QROM from the MLWE problem, derived using our new framework from Figure 2. Dilithium-QROM is essentially a less compact variant ($\approx 3X$ larger) of Dilithium with modified parameters to allow the underlying identification scheme to admit a lossy mode. We additionally prove the security of the original Dilithium scheme in the QROM based on MLWE and another non-interactive assumption.

1.1.1 Security of Fiat-Shamir Signatures

Security of deterministic Fiat-Shamir signatures $\text{DFS}[\text{ID}, \text{H}, \text{PRF}]$ in the QROM is proved in two independent steps, see Figure 2.

STEP 1: LOSSY \implies UF-NMA. We sketch an adaptation of the standard history-free proof implicitly contained in [AFLT12]. By the security properties of the lossy identification scheme, the public key can be set in lossy mode which remains unnoticed by a computationally bounded quantum adversary.




Figure 2: Security of standard Fiat-Shamir signatures $\text{FS} = \text{FS}[\text{ID}, \text{H}]$ and deterministic Fiat-Shamir signatures $\text{DFS} = \text{DFS}[\text{ID}, \text{H}, \text{PRF}]$ in the QROM. Solid arrows denote tight reductions, dashed arrows non-tight reductions. The considered security notions are: UF-CMA (unforgeability against chosen-message attack), UF-CMA₁ (unforgeability against one-query-per-message chosen-message attack), and UF-NMA (unforgeability against no-message attack).

Further, breaking the signature scheme in lossy mode with at most Q_H queries to the quantum random oracle essentially requires to solve the generic quantum search problem, whose complexity is $\Theta(Q_H^2 \cdot \varepsilon_{\text{ls}})$ [HRS16, Zha12a], where ε_{ls} is the statistical soundness parameter of ID in lossy mode. A similar argument is implicitly contained in [Unr17a, ABB⁺17].

STEP 2: UF-NMA \implies UF-CMA. We will now sketch a history-free proof of $\text{UF-NMA} \implies \text{UF-CMA}_1$, where (compared to UF-CMA security) UF-CMA₁ security limits the number of queried signatures per message M to one. We then apply a standard (history-free, tight) reduction to show that UF-CMA₁ secure signatures de-randomized with a PRF yield UF-CMA secure signatures with deterministic signing [BPS16].

The standard ROM proof of $\text{UF-NMA} \implies \text{UF-CMA}$ (implicitly contained in [AFLT12]) works as follows: one uses the HVZK property of ID to show that the signing oracle can be efficiently simulated only knowing the public-key. Concretely, the HVZK simulator generates a transcript (W, c, Z) and later “patches” the random oracle by defining $\text{H}(W \parallel M) := c$ to make (W, Z) a valid signature. The problem is that the random oracle patching (i.e., defining $\text{H}(W \parallel M) := c$) can only be done *after* the signing query on M because only then W and c are known. This renders the AFLT standard reduction non history-free. In our history-free $\text{UF-NMA} \implies \text{UF-CMA}_1$ proof, we resolve this problem as follows. We use the HVZK property to generate the transcript (W_M, c_M, Z_M) *deterministically* using message-dependent randomness. Hence, for each message M , the transcript (W_M, c_M, Z_M) is unique and can be computed at any time. This uniqueness allows us to patch the random oracle $\text{H}(W \parallel M)$ to c_M at any time of the proof (i.e., iff $W = W_M$), even before the adversary has established a signing query on message M . This trick makes the proof history-free, see Theorem 3.2. Clearly, this only works if the adversary receives at most one signature for each messages M , which is guaranteed by the UF-CMA₁ experiment.

In order to deal with (full) UF-CMA security of probabilistic Fiat-Shamir signatures $\text{FS}[\text{ID}, \text{H}]$, the above trick can be adapted to also obtain a history-free reduction, see Theorem 3.3. However, the proof is less tight as the reduction suffers from a quadratic blow-up in its running time.

Our results furthermore prove *strong* unforgeability if the identification scheme satisfies an additional property called computational unique response (CUR). CUR essentially says that it is hard to come up with two accepting transcripts with the same commitment and challenge but different responses.

1.1.2 Dilithium-QROM: A signature scheme with provable security in the QROM

The digital signature scheme Dilithium [DLL⁺17] is constructed from a canonical identification scheme using the Fiat-Shamir with aborts approach [Lyu09]. In the ROM, its security is based (via non-tight reductions) on the hardness of the MSIS and MLWE problems. We show that by increasing the size of the modulus and the dimension of the public key matrix, the resulting identification scheme admits a lossy mode such that distinguishing real from lossy keys is based on the hardness of MLWE. We can then apply our main reduction to conclude that the resulting digital signature scheme is based on the hardness of the MLWE problem.

In order to construct an identification scheme with a lossy mode, in addition to increasing the size of the modulus and the overall dimension, we also choose our prime modulus q so that the underlying ring $\mathbb{Z}_q[X]/(X^n + 1)$ has the property that all elements with coefficients less than $\sqrt{q/2}$ have an inverse

[LN17] – having all small elements be invertible is crucial to having lossiness.¹ For the same security levels as Dilithium, the total size of the public key and signature is increased by a factor of a little over 3.

1.1.3 Revisiting the Security of Dilithium

Due to the way the parameters are set, the underlying identification scheme of the original Dilithium scheme does not have a lossy mode, and so we cannot apply Theorem 3.4 in the reduction sequence in Figure 2. Nevertheless, the reduction from Theorem 3.2 is still applicable. In the classical ROM, one then obtains a reduction from MSIS to the UF-NMA scheme via the forking lemma (see Figure 1).

The main downside of this last step is that the reduction is inherently non-tight. In practice, however, parameters are set based on the hardness of the underlying MSIS problem and the non-tightness of the reduction is ignored. This is not just the case in lattice-based schemes, but is the prevalent practice for every signature scheme built via the Fiat-Shamir transform. The implicit assumption is, therefore, that the UF-NMA scheme is *exactly* as secure as MSIS (assuming that H is secure). We point out that the assumption that the UF-NMA scheme is secure is a non-interactive assumption that is reasonably simple to state, and so the fact that several decades of cryptanalysis haven’t produced any improved attacks against schemes whose parameters ignore the non-tightness of the reduction, gives us confidence that equating the hardness of the UF-NMA scheme with the hardness of the underlying problem is very reasonable.

In Section 4.5, we formulate the security of the UF-NMA scheme as a “convolution” of a lattice/hash function problem, which we call `SelfTargetMSIS`, and then show that based on the hardness of `MLWE` and `SelfTargetMSIS`, the deterministic version of the Dilithium scheme is (tightly) UF-CMA secure in the QROM. In other words, we show that the security of the *tight* version of the signature scheme is based on exactly the same assumptions in the ROM and the QROM.

1.1.4 Other Instantiations

Our framework can be applied to obtain a security proof in the QROM for a number of existing Fiat-Shamir signature schemes that are similar to Dilithium (e.g., [Lyu09, AFLT12, GLP12, Lyu12, BG14, Lyu16, ABB⁺17]) and those that have a somewhat different structure and possibly based on different assumptions (e.g., [KTX08, SSH11, DDL13]). Our rationale for setting the parameters in Dilithium-QROM was to minimize the total sum of the public key and the signature. If one, on the other hand, wished to only minimize the signature size, one could create a public key whose “height” is larger than its “width” (e.g., as in [ABB⁺17]). For optimal efficiency, this may possibly require working over polynomial rings $\mathbb{Z}_q[X]/(f(x))$ which are finite fields.

1.2 Concrete Bounds and Comparison with Unruh [Unr17a, Unr17b]

Ignoring all constants and the computational term accounting for the pseudo-random function, our concrete bound for the UF-CMA security of deterministic Fiat-Shamir signatures DFS in the QROM is

$$\text{Adv}_{\text{DFS}}^{\text{UF-CMA}}(\mathbf{A}) \leq \text{Adv}_{\text{ID}}^{\text{LOSS}}(\mathbf{B}) + Q_H^2 \cdot \varepsilon_{\text{ls}} + Q_S \cdot \varepsilon_{\text{zk}} + 2^{-\alpha}, \quad \text{Time}(\mathbf{B}) \approx \text{Time}(\mathbf{A}) \quad (1)$$

where $\text{Adv}_{\text{ID}}^{\text{LOSS}}(\mathbf{B})$ is the lossyness advantage of ID, ε_{ls} is the statistical soundness parameter of ID in lossy mode, α is the min-entropy of ID’s commitments, and ε_{zk} is the HVZK parameter of ID.

From Unruh [Unr17b] one can derive the following concrete bound which even holds for (standard) probabilistic Fiat-Shamir signatures FS.

$$\begin{aligned} \text{Adv}_{\text{FS}}^{\text{UF-CMA}}(\mathbf{A}) &\leq \text{Adv}_{\text{ID}}^{\text{LOSS}}(\mathbf{B}) + Q_H^2 \cdot \varepsilon_{\text{ls}} + Q_S \cdot \varepsilon_{\text{zk}} + Q_S Q_H^{1/2} \cdot 2^{-\alpha/4}, \\ \text{Time}(\mathbf{B}) &\approx \text{Time}(\mathbf{A}) + Q_H Q_S. \end{aligned} \quad (2)$$

Compared to (1), bound (2) has two sources of non-tightness.

The first source of non-tightness in (2) is the term $Q_S Q_H^{1/2} \cdot 2^{-\alpha/4}$ which stems from a generic re-programming technique from [Unr15]. In most practical lattice-based schemes the commitment’s min-entropy α is large enough not to make a big impact on the worse bounds. However, this term puts a

¹There do not exist q for which $\mathbb{Z}_q[X]/(X^n + 1)$ is a field.

lower bound on the min-entropy of commitments which translates to an unnatural lower bound on the size of quantum-resistant Fiat-Shamir signatures. Furthermore, it is sometimes not that easy to exactly compute the min-entropy α . Further, simple techniques to get a “good-enough” bound (as we did for regular Dilithium when we obtained $\alpha = 255$) would no longer result in something meaningful when used with (2).

The second and more important sources of non-tightness in (2) is the quadratic (in the number of queries) blow-up in the running time $\text{Time}(\mathbf{B}) \approx \text{Time}(\mathbf{A}) + Q_H Q_S$ which renders the reduction non-tight in all practical aspects. Interestingly, our proof for the security of *probabilistic* Fiat-Shamir signatures (Theorem 3.3) introduces the same source of non-tightness. However, under the assumption that superposition queries to classical data can be performed in a single time step (denoted by QRAM in [Unr17b]), the running time in (2) drops to $\text{Time}(\mathbf{B}) \approx \text{Time}(\mathbf{A})$ and hence the reduction is tight again. We leave it as an open problem to come up with a tight reduction for probabilistic Fiat-Shamir signatures in the QROM without using QRAM.

2 Preliminaries

For $n \in \mathbb{N}$, let $[n] := \{1, \dots, n\}$. For a set S , $|S|$ denotes the cardinality of S . For a finite set S , we denote the sampling of a uniform random element x by $x \leftarrow S$, while we denote the sampling according to some distribution \mathfrak{D} by $x \leftarrow \mathfrak{D}$. By $\llbracket B \rrbracket$ we denote the bit that is 1 if the Boolean Statement B is true, and 0 otherwise.

ALGORITHMS. Let \mathbf{A} be an algorithm. Unless stated otherwise, we assume all our algorithms to be probabilistic. We denote by $y \leftarrow \mathbf{A}(x)$ the probabilistic computation of algorithm \mathbf{A} on input x . If \mathbf{A} is deterministic, we write $y := \mathbf{A}(x)$. The notation $y \in \mathbf{A}(x)$ is used to indicate all possible outcomes y of the probabilistic algorithm \mathbf{A} on input x . We can make any probabilistic \mathbf{A} deterministic by running it with fixed randomness. We write $y := \mathbf{A}(x; r)$ to indicate that \mathbf{A} is run on input x with randomness r . Finally, the notation $\mathbf{A}(x) \Rightarrow y$ denotes the event that \mathbf{A} on input x returns y .

GAMES. Following [Sho04, BR06], we use code-based games. We implicitly assume boolean flags to be initialized to false, numerical types to 0, sets to \emptyset , and strings to the empty string ϵ . We make the convention that a procedure terminates once it has returned an output.

2.1 Quantum Computation

QUANTUM STATES. The state of a qubit $|\phi\rangle$ is described by a two-dimensional complex vector $|\phi\rangle = \alpha|0\rangle + \beta|1\rangle$ where $\{|0\rangle, |1\rangle\}$ form an orthonormal basis of \mathbb{C}^2 and $\alpha, \beta \in \mathbb{C}$ with $|\alpha|^2 + |\beta|^2 = 1$ are called the complex *amplitudes* of $|\phi\rangle$. The qbit $|\phi\rangle$ is said to be *in superposition* if $0 < |\alpha| < 1$. A classical bit $b \in \{0, 1\}$ is naturally encoded as state $|b\rangle$ of a qubit.

The state $|\psi\rangle$ of n qubits can be expressed as $|\psi\rangle = \sum_{x \in \{0,1\}^n} \alpha_x |x\rangle \in \mathbb{C}^{2^n}$ where $\{\alpha_x\}_{x \in \{0,1\}^n}$ is a set of 2^n complex amplitudes such that $\sum_{x \in \{0,1\}^n} |\alpha_x|^2 = 1$. As for one qubit, the standard orthonormal or *computational basis* is given by $\{|x\rangle\}_{x \in \{0,1\}^n}$. When the quantum state $|\psi\rangle$ is *measured* in the computational basis, the outcome is the classical string $x \in \{0, 1\}^n$ with probability $|\alpha_x|^2$ and the quantum state collapses to what is observed, namely $|x\rangle$.

The evolution of a quantum system in state $|\psi\rangle$ can be described by a linear length-preserving transformation $U : \mathbb{C}^{2^n} \rightarrow \mathbb{C}^{2^n}$. Such transformations correspond to *unitary* matrices U of size 2^n by 2^n , i.e. U has the property that $UU^\dagger = \mathbb{1}$, where U^\dagger is the complex-conjugate transpose of U .

For further details about basic concepts and notation of quantum computing, we refer to the standard text book by Nielsen and Chuang [NC00].

QUANTUM ORACLES AND QUANTUM ADVERSARIES. For a classical oracle function $\mathbf{O} : \{0, 1\}^n \rightarrow \{0, 1\}^m$, we follow the standard approach as in [BBC⁺98, BDF⁺11] to make the execution of the classical function \mathbf{O} a reversible unitary transformation. We model quantum access to \mathbf{O} by

$$U_{\mathbf{O}} : |x\rangle|y\rangle \mapsto |x\rangle|y \oplus \mathbf{O}(x)\rangle ,$$

where $x \in \{0, 1\}^n$ and $y \in \{0, 1\}^m$. Note that due to the XOR function in the second register, $U_{\mathbf{O}}$ is

its own inverse, i.e. executing U_O twice results in the identity for any function O .² Quantum oracle adversaries $A^{(O)}$ can access O in superposition by applying U_O . The quantum time it takes to apply U_O is linear in the time it takes to evaluate O classically. We write $A^{(O)}$ to indicate that an oracle is quantum-accessible, contrary to oracles which can only be accessed classically which are denoted by A^O . We also abuse notation and use $|O\rangle$ to denote the oracle that is quantumly accessible.

QUANTUM RANDOM-ORACLE MODEL. We consider security games in the quantum random-oracle model (QROM) [BDF⁺11] like their counterparts in the classical random-oracle model [BR93], with the difference that we consider quantum adversaries that are given **quantum** access to the random oracles involved, and **classical** access to all other oracles (e.g., the signing oracle). Zhandry [Zha12b] proved that no quantum algorithm $A^{(H)}$, issuing at most Q quantum queries to $|H\rangle$, can distinguish between a random function $H : \{0, 1\}^m \rightarrow \{0, 1\}^n$ and a $2Q$ -wise independent function f_{2Q} . For concreteness, we view $f_{2Q} : \{0, 1\}^m \rightarrow \{0, 1\}^n$ as a random polynomial of degree $2Q$ over the finite field \mathbb{F}_{2^n} . The running time to evaluate f_{2Q} is linear in Q .

In this article, we will use this observation in the context of security reductions, where quantum adversary B simulates quantum adversary $A^{(H)}$ which makes at most Q queries to $|H\rangle$. Hence, the running time of B is $\text{Time}(B) = \text{Time}(A) + q \cdot \text{Time}(H)$, where $\text{Time}(H)$ is the time it takes to simulate $|H\rangle$. Using the observation above, B can use a $2Q$ -wise independent function in order to (information-theoretically) simulate $|H\rangle$ and we obtain that the running time of B is $\text{Time}(B) = \text{Time}(A) + Q \cdot \text{Time}(f_{2Q})$, and the time $\text{Time}(f_{2Q})$ to evaluate f_{2Q} is linear in Q . The second term of this running time (quadratic in Q) can be further reduced to linear in Q in the quantum random-oracle model where B can simply use another random oracle to simulate $|H\rangle$. Assuming evaluating the random oracle takes one time unit, we write $\text{Time}(B) = \text{Time}(A) + Q$ which is approximately $\text{Time}(A)$.

GENERIC QUANTUM SEARCH. For $\lambda \in [0, 1]$ let \mathcal{B}_λ be the Bernoulli distribution, i.e., $\Pr[b = 1] = \lambda$ for the bit $b \leftarrow \mathcal{B}_\lambda$. Let X be some finite set. The generic quantum search problem GSP [HRS16, Zha12a] is to find an $x \in X$ satisfying $g(x) = 1$ given quantum access to an oracle $g : X \rightarrow \{0, 1\}$, such that for each $x \in X$, $g(x)$ is distributed according to \mathcal{B}_λ . We will need the following slight variation of GSP. The Generic quantum Search Problem with Bounded probabilities GSPB is like the quantum search problem with the difference that the Bernoulli parameter $\lambda(x)$ may (adversarially) depend on x but it is upper bounded by a global λ .

Lemma 2.1 (*Generic Search Problem with Bounded Probabilities*) *Let $\lambda \in [0, 1]$. For any (unbounded, quantum) algorithm A issuing at most Q quantum queries to $|g\rangle$, $\Pr[\text{GSPB}_\lambda^A \Rightarrow 1] \leq 8 \cdot \lambda \cdot (Q + 1)^2$, where Game GSPB_λ is defined in Figure 3.*

GAME GSPB_λ
01 $(\lambda(x))_{x \in X} \leftarrow A_1$
02 If $\exists x \in X$ s.t. $\lambda(x) > \lambda$ then **return** 0
03 For all $x \in X$: $g(x) \leftarrow \mathcal{B}_{\lambda(x)}$
04 $x \leftarrow A_2^{(g)}$
05 **return** $g(x)$

Figure 3: The generic search game GSPB_λ with bounded maximal Bernoulli parameter $\lambda \in [0, 1]$.

The bound on GSPB can be reduced to the known bound on GSP [HRS16, Zha12a] by artificially increasing the Bernoulli parameter to obtain the dependence on each $x \in X$. Proof details are given in Appendix A.1.

2.2 Pseudorandom Functions

A pseudorandom function PRF is a mapping $\text{PRF} : \mathcal{K} \times \{0, 1\}^n \rightarrow \{0, 1\}^k$, where \mathcal{K} is a finite key space and n, k are integers. To a quantum adversary A and PRF we associate the advantage function

$$\text{Adv}_{\text{PRF}}^{\text{PR}}(A) := \left| \Pr[A^{\text{PRF}(K, \cdot)} \Rightarrow 1 \mid K \leftarrow \mathcal{K}] - \Pr[A^{\text{RF}(\cdot)} \Rightarrow 1] \right|,$$

²Together with the observation that taking the conjugate-complex and transposing U_O do not change U_O , we obtain $U_O^\dagger = U_O$, and hence, $U_O U_O^\dagger = U_O^2 = \mathbb{1}$, showing that U_O is indeed a unitary transformation.




Figure 4: A canonical identification scheme and its transcript (W, c, Z) .

where $\text{RF} : \{0, 1\}^n \rightarrow \{0, 1\}^k$ is a perfect random function. We note that while adversary A is quantum, it only gets classical access to the oracles $\text{PRF}(K, \cdot)$ and $\text{RF}(\cdot)$.

2.3 Canonical Identification Schemes

A canonical identification scheme ID is a three-move protocol of the form depicted in Figure 4. The prover's first message W is called *commitment*, the verifier selects a uniform *challenge* c from set ChSet , and, upon receiving a *response* Z from the prover, makes a deterministic decision.

Definition 2.2 (Canonical Identification Scheme). A canonical identification scheme ID is defined as a tuple of algorithms $\text{ID} := (\text{IGen}, P, \text{ChSet}, V)$.

- The key generation algorithm IGen takes system parameters par as input and returns public and secret key (pk, sk) . We assume that pk defines ChSet (the set of challenges), WSet (the set of commitments), and ZSet (the set of responses).
- The prover algorithm $P = (P_1, P_2)$ is split into two algorithms. P_1 takes as input the secret key sk and returns a commitment $W \in \text{WSet}$ and a state St ; P_2 takes as input the secret key sk , a commitment W , a challenge c , and a state St and returns a response $Z \in \text{ZSet} \cup \{\perp\}$, where $\perp \notin \text{ZSet}$ is a special symbol indicating failure.
- The verifier algorithm V takes the public key pk and the conversation transcript as input and outputs a *deterministic decision*, 1 (acceptance) or 0 (rejection).

We make a couple of useful definitions. A *transcript* is a three-tuple $(W, c, Z) \in \text{WSet} \times \text{ChSet} \times \text{ZSet} \cup \{\perp, \perp, \perp\}$. It is called *valid* (with respect to public-key pk) if $V(pk, W, c, Z) = 1$. In Figure 5 we also define a transcript oracle Trans that returns a real interaction (W, c, Z) between prover and verifier as depicted in Figure 4, with the important convention that the transcript is defined as (\perp, \perp, \perp) if $Z = \perp$.

<p>Algorithm $\text{Trans}(sk)$:</p> 01 $(W, St) \leftarrow P_1(sk)$ 02 $c \leftarrow \text{ChSet}$ 03 $Z \leftarrow P_2(sk, W, c, St)$ 04 if $Z = \perp$ then return (\perp, \perp, \perp) 05 return (W, c, Z)

Figure 5: An honestly generated transcript (W, c, Z) output by the transcript oracle $\text{Trans}(sk)$.

Definition 2.3 (Correctness Error). Identification scheme ID has correctness error δ if for all $(pk, sk) \in \text{IGen}(\text{par})$ the following holds:

- All possible transcripts (W, c, Z) satisfying $Z \neq \perp$ are valid, i.e., for all $(W, St) \in P_1(sk)$, all $c \in \text{ChSet}$ and all $Z \in P_2(sk, W, c, St)$ with $Z \neq \perp$, we have $V(pk, W, c, Z) = 1$.
- The probability that an honestly generated transcript (W, c, Z) contains $Z = \perp$ is bounded by δ , i.e., $\Pr[Z = \perp \mid (W, c, Z) \leftarrow \text{Trans}(sk)] \leq \delta$.

Definition 2.4 We call ID *commitment-recoverable*, if for any $(pk, sk) \in \text{IGen}(\text{par})$, $c \in \text{ChSet}$, and $Z \in \text{ZSet}$, there exists a unique $W \in \text{WSet}$ such that $V(pk, W, c, Z) = 1$. This unique W can be publicly computed using a commitment recovery algorithm as $W := \text{Rec}(pk, c, Z)$.

We define no-abort honest-verifier zero-knowledge, a weak variant of honest-verifier zero-knowledge that requires the transcript (as generated by $\text{Trans}(sk)$) to be publicly simulatable, conditioned on $Z \neq \perp$.

Definition 2.5 (No-Abort Honest-verifier Zero-knowledge). A canonical identification scheme ID is said to be ε_{zk} -perfect naHVZK (no-abort honest-verifier zero-knowledge) if there exists an algorithm Sim that, given only the public key pk , outputs (W, c, Z) such that the following conditions hold:

- The distribution of $(W, c, Z) \leftarrow \text{Sim}(pk)$ has statistical distance at most ε_{zk} from $(W', c', Z') \leftarrow \text{Trans}(sk)$, where Trans is defined in Figure 5.
- The distribution of c from $(W, c, Z) \leftarrow \text{Sim}(pk)$ conditioned on $c \neq \perp$ is uniform random in ChSet .

Note that if ID is commitment-recoverable, then we can abandon the W in the output of Trans and Sim since W can be publicly computed from (c, Z) .

Definition 2.6 (Min-Entropy). If the most likely value of a random variable W that is chosen from a discrete distribution D occurs with probability $2^{-\alpha}$, then we say that $\text{min-entropy}(W \mid W \leftarrow D) = \alpha$. We will say that a canonical identification scheme ID has α bits of min-entropy, if

$$\Pr_{(pk, sk) \leftarrow \text{IGen}(\text{par})} [\text{min-entropy}(W \mid (W, St) \leftarrow P_1(sk)) \geq \alpha] \geq 1 - 2^{-\alpha}.$$

In other words, except with probability $2^{-\alpha}$ over the choice of (pk, sk) , the min-entropy of W will be at least α .

An identification scheme has unique responses if for all W and c there exists at most one Z to make the verifier accept, i.e., $V(pk, W, c, Z) = 1$. We relax this property to computational unique response (CUR) for which we require it to be computationally difficult to come up with (W, c, Z, Z') with $V(pk, W, c, Z) = V(pk, W, c, Z') = 1$ and $Z' \neq Z$.

Definition 2.7 (Computational Unique Response). To an adversary A we associate the advantage function

$$\text{Adv}_{\text{ID}}^{\text{CUR}}(A) := \Pr \left[\begin{array}{l} V(pk, W, c, Z) = 1 \\ V(pk, W, c, Z') = 1 \wedge Z \neq Z' \end{array} \mid \begin{array}{l} (pk, sk) \leftarrow \text{IGen}(\text{par}); \\ (W, c, Z, Z') \leftarrow A(pk) \end{array} \right].$$

LOSSY IDENTIFICATION SCHEMES. We now recall lossy identification schemes [AFLT12].

Definition 2.8 An identification scheme $\text{ID} = (\text{IGen}, P, \text{ChSet}, V)$ is lossy if there exists a lossy key generation algorithm LossyGen that takes system parameters par as input and returns public key pk_{ls} (and no secret key sk).

We refer to $\text{LID} = (\text{IGen}, \text{LossyGen}, P, \text{ChSet}, V)$ as a lossy identification scheme.

We now define two security properties of a lossy identification scheme LID. The first property says that public keys generated with the real key generator IGen are indistinguishable from ones generated by the lossy key generator LossyGen . Concretely, we define the *LOSS advantage function of a quantum adversary A against ID* as

$$\text{Adv}_{\text{LID}}^{\text{LOSS}}(A) := \left| \Pr[A(pk_{\text{ls}}) \Rightarrow 1 \mid pk_{\text{ls}} \leftarrow \text{LossyGen}(\text{par})] - \Pr[A(pk) \Rightarrow 1 \mid (pk, sk) \leftarrow \text{IGen}(\text{par})] \right|.$$

The second security property is statistical and says that relative to a lossy key pk_{ls} , not even an unbounded quantum adversary can impersonate the prover. We say that ID has ε_{ls} -lossy soundness if for every (possibly unbounded, quantum) adversary C , $\Pr[\text{LOSSY-IMP}^C \Rightarrow 1] \leq \varepsilon_{\text{ls}}$, where game LOSSY-IMP is defined in Figure 6.

Since C is unbounded, we can upper bound $\Pr[\text{LOSSY-IMP}^C \Rightarrow 1]$ as

$$\Pr[\text{LOSSY-IMP}^C \Rightarrow 1] \leq \mathbf{E} \left[\max_{W \in \text{WSet}} \left(\Pr_{c \leftarrow \text{ChSet}} [\exists Z \in \text{ZSet} : V(pk_{\text{ls}}, W, c, Z) = 1] \right) \right), \quad (3)$$

where the expectation is taken over $pk_{\text{ls}} \leftarrow \text{LossyGen}(\text{par})$. Note that equality in Equation (3) is achieved for the “optimal” adversary C which on the “easiest” commitment $W \in \text{WSet}$ and a random challenge $c \leftarrow \text{ChSet}$ finds a response $Z \in \text{ZSet}$ that the verifier accepts.

GAME LOSSY-IMP: 01 $pk_{\text{ls}} \leftarrow \text{LossyGen}(\text{par})$ 02 $(W^*, St) \leftarrow C(pk_{\text{ls}})$ 03 $c^* \leftarrow \text{ChSet}$ 04 $Z^* \leftarrow C(St, c^*)$ 05 return $\llbracket V(pk_{\text{ls}}, W^*, c^*, Z^*) \rrbracket$
--

Figure 6: The lossy impersonation game LOSSY-IMP.

2.4 Digital Signatures

We now define syntax and security of a digital signature scheme. Let par be common system parameters shared among all participants.

Definition 2.9 (Digital Signature). A digital signature scheme SIG is defined as a triple of algorithms $\text{SIG} = (\text{Gen}, \text{Sign}, \text{Ver})$.

- The key generation algorithm $\text{Gen}(\text{par})$ returns the public and secret keys (pk, sk) . We assume that pk defines the message space MSet .
- The signing algorithm $\text{Sign}(sk, M)$ returns a signature σ .
- The deterministic verification algorithm $\text{Ver}(pk, M, \sigma)$ returns 1 (accept) or 0 (reject).

Signature scheme SIG has correctness error γ if for all $(pk, sk) \in \text{Gen}(\text{par})$, all messages $M \in \text{MSet}$, we have $\Pr[\text{Ver}(pk, M, \text{Sign}(sk, M)) = 0] \leq \gamma$.

SECURITY. We define the UF-CMA (unforgeability against chosen-message attack), UF-CMA₁ (unforgeability against one-per-message chosen-message attack), and UF-NMA (unforgeability against no-message attack) advantage functions of a quantum adversary A against SIG as $\text{Adv}_{\text{SIG}}^{\text{UF-CMA}}(A) := \Pr[\text{UF-CMA}^A \Rightarrow 1]$, $\text{Adv}_{\text{SIG}}^{\text{UF-CMA}_1}(A) := \Pr[\text{UF-CMA}_1^A \Rightarrow 1]$, and $\text{Adv}_{\text{SIG}}^{\text{UF-NMA}}(A) := \Pr[\text{UF-NMA}^A \Rightarrow 1]$, where the games UF-CMA, UF-CMA₁, and UF-NMA are given in Figure 7. We also consider *strong* unforgeability where the adversary may return a forgery on a message previously queried to the signing oracle, but with a different signature. In the corresponding experiments sUF-CMA and sUF-CMA₁, the set \mathcal{M} contains tuples (M, σ) and for the winning condition it is checked that $(M^*, \sigma^*) \notin \mathcal{M}$.

GAMES UF-CMA/UF-CMA₁/UF-NMA:	SIGN(M)	SIGN₁(M)
01 $(pk, sk) \leftarrow \text{Gen}(\text{par})$	06 $\mathcal{M} = \mathcal{M} \cup \{M\}$	09 if $M \in \mathcal{M}$ then return \perp
02 $(M^*, \sigma^*) \leftarrow A^{\text{SIGN}}(pk)$ //UF-CMA	07 $\sigma \leftarrow \text{Sign}(sk, M)$	10 $\mathcal{M} = \mathcal{M} \cup \{M\}$
03 $(M^*, \sigma^*) \leftarrow A^{\text{SIGN}_1}(pk)$ //UF-CMA ₁	08 return σ	11 $\sigma \leftarrow \text{Sign}(sk, M)$
04 $(M^*, \sigma^*) \leftarrow A(pk)$ //UF-NMA		12 return σ
05 return $\llbracket M^* \notin \mathcal{M} \rrbracket \wedge \text{Ver}(pk, M^*, \sigma^*)$		

Figure 7: Games UF-CMA, UF-CMA₁, and UF-NMA.

Any UF-CMA₁ (sUF-CMA₁) secure signature scheme can be combined with a pseudo-random function PRF to obtain an UF-CMA (sUF-CMA) secure signature scheme by defining $\text{Sign}'((sk, K), M) := \text{Sign}(sk, M; \text{PRF}_K(M))$, where K is a secret PRF key which is part of the secret key. This construction is well known in the classical setting [BPS16], and the same proof works in the quantum setting. Here PRF only has to provide security against quantum adversaries where the access to PRF is classical.

3 Fiat-Shamir in the Quantum Random-Oracle Model

3.1 Signatures from Identification Schemes

Let $\text{ID} := (\text{IGen}, \text{P}, \text{ChSet}, \text{V})$ be a canonical identification scheme, let κ_m be a positive integer, and let $\text{H} : \{0, 1\}^* \rightarrow \text{ChSet}$ be a hash function. The following signature scheme $\text{SIG} := (\text{Gen} = \text{IGen}, \text{Sign}, \text{Ver})$ is obtained by the Fiat-Shamir transformation with aborts $\text{FS}[\text{ID}, \text{H}, \kappa_m]$ [Lyu09].

$\text{Sign}(sk, M)$	$\text{Ver}(pk, M, \sigma)$
01 $\kappa := 0$	09 Parse $\sigma = (W, Z) \in \text{WSet} \times \text{ZSet}$
02 while $Z = \perp$ and $\kappa \leq \kappa_m$ do	10 $c = \text{H}(W \parallel M)$
03 $\kappa := \kappa + 1$	11 return $\mathbb{V}(pk, W, c, Z) \in \{0, 1\}$
04 $(W, St) \leftarrow \text{P}_1(sk)$	
05 $c = \text{H}(W \parallel M)$	
06 $Z \leftarrow \text{P}_2(sk, W, c, St)$	
07 if $Z = \perp$ return $\sigma = \perp$	
08 return $\sigma = (W, Z)$	

We make the convention that if $\sigma = (W, Z)$ is not in $\text{WSet} \times \text{ZSet}$, then $\text{Ver}(pk, M, \sigma)$ returns 0 (reject). Clearly, if ID has correctness error δ , then SIG has correctness error $\gamma = \delta^{\kappa_m}$.

FIAT-SHAMIR FOR COMMITMENT-RECOVERABLE IDENTIFICATION. For commitment-recoverable ID (see Definition 2.4), we can define an alternative Fiat-Shamir transformation $\text{SIG}' = \text{FS}'[\text{ID}, \text{H}, \kappa_m] := (\text{Gen} = \text{IGen}, \text{Sign}', \text{Ver}')$. Algorithm $\text{Sign}'(sk, M)$ is defined as $\text{Sign}(sk, M)$ with the modified output $\sigma' = (c, Z)$. Algorithm $\text{Ver}'(pk, M, \sigma')$ first parses $\sigma' = (c, Z)$, then recomputes the commitment as $W' := \text{Rec}(pk, c, Z)$, and finally returns 1 iff $\text{H}(W' \parallel M) = c$.

$\text{Sign}'(sk, M)$	$\text{Ver}'(pk, M, \sigma')$
01 $\kappa := 0$	09 Parse $\sigma' = (c, Z) \in \text{ChSet} \times \text{ZSet}$
02 while $Z = \perp$ and $\kappa \leq \kappa_m$ do	10 $W' := \text{Rec}(pk, c, Z)$
03 $\kappa := \kappa + 1$	11 return $\llbracket \text{H}(W' \parallel M) = c \rrbracket$
04 $(W, St) \leftarrow \text{P}_1(sk)$	
05 $c = \text{H}(W \parallel M)$	
06 $Z \leftarrow \text{P}_2(sk, W, c, St)$	
07 if $Z = \perp$ return $\sigma' = \perp$	
08 return $\sigma' = (c, Z)$	

Since $\sigma = (W, Z)$ can be publicly transformed into $\sigma' = (c, Z)$ and vice versa, SIG and SIG' are equivalent in terms of security. The alternative Fiat-Shamir transform yields shorter signatures if $c \in \text{ChSet}$ has a smaller representation size than the commitment $W \in \text{WSet}$.

MAIN SECURITY STATEMENT. The following is our main security statement for $\text{SIG} := \text{FS}[\text{ID}, \text{H}, \kappa_m]$ in the QROM.

Theorem 3.1 *Assume the identification scheme ID is lossy, ε_{zk} -perfect naHVZK, has α bits of min entropy, and is ε_{ls} -lossy sound. For any quantum adversary A against UF-CMA₁ (sUF-CMA₁) security that issues at most Q_{H} queries to the quantum random oracle |H⟩ and Q_{S} classical queries to the signing oracle SIGN₁, there exists a quantum adversary B (and a quantum adversary C against CUR) such that*

$$\begin{aligned} \text{Adv}_{\text{SIG}}^{\text{UF-CMA}_1}(\text{A}) &\leq \text{Adv}_{\text{ID}}^{\text{LOSS}}(\text{B}) + 8(Q_{\text{H}} + 1)^2 \cdot \varepsilon_{\text{ls}} + \kappa_m Q_{\text{S}} \cdot \varepsilon_{\text{zk}} + 2^{-\alpha+1} , \\ \text{Adv}_{\text{SIG}}^{\text{sUF-CMA}_1}(\text{A}) &\leq \text{Adv}_{\text{ID}}^{\text{LOSS}}(\text{B}) + 8(Q_{\text{H}} + 1)^2 \cdot \varepsilon_{\text{ls}} + \kappa_m Q_{\text{S}} \cdot \varepsilon_{\text{zk}} + 2^{-\alpha+1} + \text{Adv}_{\text{ID}}^{\text{CUR}}(\text{C}) , \end{aligned}$$

and $\text{Time}(\text{B}) = \text{Time}(\text{C}) = \text{Time}(\text{A}) + \kappa_m Q_{\text{H}} \approx \text{Time}(\text{A})$.

Note that with this observation the bound of Theorem 3.1 is tight, i.e., the computational advantages appear with a constant factor (one). In the classical ROM setting, the only difference is that the bound depends linearly on Q_{H} , instead of quadratic.

DETERMINISTIC FIAT-SHAMIR. Let PRF be a pseudo-random function. Consider a deterministic variant $\text{DSIG} := \text{DFS}[\text{ID}, \text{H}, \text{PRF}, \kappa_m] = (\text{Gen}, \text{DSign}, \text{Ver})$ of FS where lines 04 and 06 of Sign is derandomized using the PRF, where the random key K is part of the secret key.

$\text{DSign}((sk, K), M)$
01 $\kappa := 0$
02 while $Z = \perp$ and $\kappa \leq \kappa_m$ do
03 $\kappa := \kappa + 1$
04 $(W, St) := \text{P}_1(sk; \text{PRF}_K(0 \parallel m \parallel \kappa))$
05 $c = \text{H}(W \parallel M)$
06 $Z := \text{P}_2(sk, W, c, St; \text{PRF}_K(1 \parallel m \parallel \kappa))$
07 if $Z = \perp$ return $\sigma = \perp$
08 return $\sigma = (W, Z)$

As discussed at the end of Section 2.4, the UF-CMA (sUF-CMA) security of DSIG is implied by the UF-CMA₁ (sUF-CMA₁) security of FS. Concretely the advantages are upper bounded by the same terms as in Theorem 3.1 plus an additional term $\text{Adv}_{\text{PRF}}^{\text{PR}}(\text{D})$ accounting for the quantum security of the PRF.

$$\begin{aligned}\text{Adv}_{\text{DSIG}}^{\text{UF-CMA}}(\text{A}) &\leq \text{Adv}_{\text{ID}}^{\text{LOSS}}(\text{B}) + 8 \cdot (Q_{\text{H}} + 1)^2 \cdot \varepsilon_{\text{ls}} + \text{Adv}_{\text{PRF}}^{\text{PR}}(\text{D}) + \kappa_m Q_S \cdot \varepsilon_{\text{zk}} + 2^{-\alpha+1} \\ \text{Adv}_{\text{DSIG}}^{\text{sUF-CMA}}(\text{A}) &\leq \text{Adv}_{\text{ID}}^{\text{LOSS}}(\text{B}) + 8 \cdot (Q_{\text{H}} + 1)^2 \cdot \varepsilon_{\text{ls}} + \text{Adv}_{\text{PRF}}^{\text{PR}}(\text{D}) + \kappa_m Q_S \cdot \varepsilon_{\text{zk}} + 2^{-\alpha+1} + \text{Adv}_{\text{ID}}^{\text{CUR}}(\text{C})\end{aligned}$$

3.2 Security proof

The proof of Theorem 3.1 is modular. First, in Theorem 3.2 we prove that UF-NMA security plus naHVZK implies UF-CMA₁ security. Second, in Theorem 3.4 we prove that a lossy identification scheme is always UF-NMA secure.

3.2.1 From UF-NMA to UF-CMA₁

Theorem 3.2 *Assume the identification scheme ID is ε_{zk} -perfect naHVZK and has α bits of min entropy. For any UF-CMA₁ (sUF-CMA₁) quantum adversary A that issues at most Q_{H} queries to the quantum random oracle |H) and Q_S (classical) queries to the signing oracle SIGN₁, there exists a quantum adversary B against UF-NMA security making Q_{H} queries to its own quantum random oracle (and a quantum adversary C against CUR) such that*

$$\begin{aligned}\text{Adv}_{\text{SIG}}^{\text{UF-CMA}_1}(\text{A}) &\leq \text{Adv}_{\text{SIG}}^{\text{UF-NMA}}(\text{B}) + 2^{-\alpha+1} + \kappa_m Q_S \cdot \varepsilon_{\text{zk}} \\ \text{Adv}_{\text{SIG}}^{\text{sUF-CMA}_1}(\text{A}) &\leq \text{Adv}_{\text{SIG}}^{\text{UF-NMA}}(\text{B}) + 2^{-\alpha+1} + \text{Adv}_{\text{ID}}^{\text{CUR}}(\text{C}) + \kappa_m Q_S \cdot \varepsilon_{\text{zk}} ,\end{aligned}$$

and $\text{Time}(\text{B}) = \text{Time}(\text{C}) = \text{Time}(\text{A}) + \kappa_m(Q_{\text{H}} + Q_S) \approx \text{Time}(\text{A})$.

Proof of Theorem 3.2. We first prove standard unforgeability (UF-CMA₁ security) and then show how the proof can be modified to obtain strong unforgeability (sUF-CMA₁ security). Let A be a quantum adversary against the UF-CMA₁ security of SIG, issuing at most Q_{H} queries to |H) and at most Q_S queries to SIGN₁. Consider the games given in Figure 8. Recall that A has classical access to the signing oracle SIGN₁ and quantum access to the random oracle H. The quantum random oracle H is called with $|W \parallel M\rangle$ and returns $|H(|W \parallel M)\rangle$. The games in Figure 8 describe the computation that is performed for any $|W \parallel M\rangle$ that has a non-zero amplitude in $|W \parallel M\rangle$.

<p>GAME G_0-G_2</p> <p>01 $(pk, sk) \leftarrow \text{IGen}(\text{par})$</p> <p>02 $(M^*, \sigma^*) \leftarrow \text{A}^{ \text{H}\rangle, \text{SIGN}_1}(pk)$</p> <p>03 Parse $\sigma^* = (W^*, Z^*)$</p> <p>04 $c^* := \text{H}(W^* \parallel M^*)$</p> <p>05 if $c^* \neq \text{H}'(W^* \parallel M^*)$ then return 0</p> <p>06 return $\llbracket M^* \notin \mathcal{M} \rrbracket \wedge \vee(pk, W^*, c^*, Z^*)$</p> <p>GetTrans($M$)</p> <p>07 $\kappa := 0$</p> <p>08 while $Z_M = \perp$ and $\kappa \leq \kappa_m$ do</p> <p>09 $\kappa := \kappa + 1$</p> <p>10 $(W_M, St) := \text{P}_1(sk; \text{RF}(0 \parallel M \parallel \kappa))$</p> <p>11 $c_M := \text{H}(W_M \parallel M)$</p> <p>12 $Z_M := \text{P}_2(sk, W_M, c_M, St; \text{RF}(1 \parallel M \parallel \kappa))$</p> <p>13 if $Z_M = \perp$ then $(W_M, c_M, Z_M) = (\perp, \perp, \perp)$</p> <p>14 return (W_M, c_M, Z_M)</p>	<p>SIGN₁(M)</p> <p>15 if $M \in \mathcal{M}$ then return \perp</p> <p>16 $\mathcal{M} = \mathcal{M} \cup \{M\}$</p> <p>17 $(W_M, c_M, Z_M) := \text{GetTrans}(M)$</p> <p>18 return $\sigma_M := (W_M, Z_M)$</p> <p>$\text{H}(W \parallel M)$ //quantum access</p> <p>19 $(W_M, c_M, Z_M) := \text{GetTrans}(M)$ // G_1-G_2</p> <p>20 if $W = W_M$ then return $c := c_M$ // G_1-G_2</p> <p>21 return $c := \text{H}'(W \parallel M)$</p> <p>GetTrans($M$) // G_1-G_2</p> <p>22 $\kappa := 0$</p> <p>23 while $Z_M = \perp$ and $\kappa \leq \kappa_m$ do</p> <p>24 $\kappa := \kappa + 1$</p> <p>25 $(W_M, c_M, Z_M) := \text{Sim}(pk; \text{RF}(M \parallel \kappa))$</p> <p>26 if $Z_M = \perp$ then $(W_M, c_M, Z_M) = (\perp, \perp, \perp)$</p> <p>27 return (W_M, c_M, Z_M)</p>
---	---

Figure 8: Games G_0, G_1, G_2 for the proof of Theorem 3.2. Here RF and H' are perfect random function that cannot be accessed by A. Deterministic algorithm GetTrans(M) is only used internally and cannot be accessed by A.

GAME G_0 . Note that game G_0 is the original UF-CMA₁ game. The signing oracle SIGN₁ produces a signature using internal deterministic algorithm GetTrans which, in lines 10 and 12, derives the randomness of P₁ and P₂ using a perfect random function RF that cannot be accessed by A. Since in the UF-CMA₁ game only one single signing query is allowed per message,

$$\Pr[G_0^A \Rightarrow 1] = \text{Adv}_{\text{SIG}}^{\text{UF-CMA}_1}(A) .$$

GAME G_1 . This game computes the signatures on M using the naHVZK simulation algorithm Sim and patches the quantum random oracle H accordingly.

Concretely, consider a classical query SIGN₁(M) and let κ_M be the smallest integer $1 \leq \kappa \leq \kappa_m$ satisfying $(W, c, Z) := \text{Sim}(pk; \text{RF}(M \parallel \kappa))$ and $Z \neq \perp$. If no such integer exists, then we define $\kappa_M := \perp$. It deterministically computes

$$(W_M, c_M, Z_M) := \text{GetTrans}(M) = \begin{cases} \text{Sim}(pk; \text{RF}(M \parallel \kappa_M)) & 1 \leq \kappa_M \leq \kappa_m \\ (\perp, \perp, \perp) & \kappa_M = \perp \end{cases} \quad (4)$$

The signature on M is returned as

$$\sigma_M := (W_M, Z_M).$$

By the naHVZK property and the union bound, the distribution of each σ_M has statistical distance at most $\kappa_m \varepsilon_{zk}$ from one computed in game G_0 . To ensure that σ_M is a valid signature on M , in line 20 the random oracle is patched such that $\text{H}(W_M \parallel M) = c_M$ holds. Concretely, a query $W \parallel M$ to quantum random oracle H with non-zero amplitude is patched with $\text{H}(W \parallel M) := c_M$ iff $W = W_M$, where c_M and W_M are computed by GetTrans(M), see Equation (4). Note that the output distribution of the random oracle H in this game remains unchanged since c_M generated by the naHVZK simulator Sim is required to be uniformly distributed.

Overall, by a union bound we obtain

$$|\Pr[G_1^A \Rightarrow 1] - \Pr[G_0^A \Rightarrow 1]| \leq \kappa_m Q_S \cdot \varepsilon_{zk} .$$

GAME G_2 . This game returns 0 in line 05 if $c^* \neq \text{H}'(W^* \parallel M^*)$. Games G_1 and G_2 can only differ if $W_{M^*} = W^*$ and $M^* \notin \mathcal{M}$. (In that case G_2 returns 0 and G_1 returns 1.) Since $M^* \notin \mathcal{M}$, the random variable W_{M^*} was not yet revealed as part of an established signature and is completely hidden from the view of the adversary. It has α bits of min-entropy, meaning we have $\Pr[W_{M^*} = W^*] \leq 2^{-\alpha}$. We obtain

$$|\Pr[G_2^A \Rightarrow 1] - \Pr[G_1^A \Rightarrow 1]| \leq 2^{-\alpha+1} .$$

Consider adversary B against the UF-NMA game from Figure 9 having quantum access to random oracle H'. It perfectly simulates A's view in game G_2 , using its own random oracle H' to simulate H' and perfectly simulating the random function RF with a $2\kappa_m Q_H$ -wise independent hash function. Assume A's forgery (M^*, σ^*) is valid in game G_2 , i.e., $M^* \notin \mathcal{M}$ and $\text{V}(pk, W^*, c^*, Z^*) = 1$, where $c^* = \text{H}(W^* \parallel M^*)$. If $\text{H}(W^* \parallel M^*) = \text{H}'(W^* \parallel M^*)$, then (M^*, σ^*) is also a valid forgery in the UF-NMA game, i.e., $\text{V}(pk, W^*, c^*, Z^*) = 1$, where $c^* = \text{H}'(W^* \parallel M^*)$. Hence,

$$\Pr[G_2^A \Rightarrow 1] = \text{Adv}_{\text{SIG}}^{\text{UF-NMA}}(\text{B}) .$$

The proof of UF-CMA₁ security follows by collecting the probabilities. The running time Time(B) of adversary B is given by the time Time(A) to run A as a blackbox in game G_2 where in every of the Q_H oracle- and Q_S signature-queries, at most $O(\kappa_m)$ computations need to be performed.

STRONG UNFORGEABILITY. For sUF-CMA₁ security we consider exactly the same games with the difference that in all games the winning condition in line 06 is changed to $\llbracket (M^*, \sigma^*) \notin \mathcal{M} \rrbracket \wedge \text{V}(pk, W^*, c^*, Z^*)$ to account for strong unforgeability, where \mathcal{M} now records all tuples (M, σ_M) of previously established messages/signature pairs.

The difference between games G_1 and G_2 is that game G_2 returns 0 in line 05 if $c^* \neq \text{H}'(W^* \parallel M^*)$, i.e., if $\text{H}(W^* \parallel M^*)$ was previously patched in line 20 with $\text{H}(W^* \parallel M^*) := c_{M^*}$. Games G_1 and G_2 can only differ if $W_{M^*} = W^*$, $(M^*, \sigma^*) \notin \mathcal{M}$, and $\text{V}(pk, W^*, c^*, Z^*) = 1$. (In that case G_2 returns 0 and G_1 returns 1.)

<p>Adversary $B^{ H'}(pk)$</p> <p>01 $(M^*, \sigma^*) \leftarrow A^{ H}, \text{SIGN}_1(pk)$</p> <p>02 Parse $\sigma^* = (W^*, Z^*)$</p> <p>03 $c^* := H(W^* \parallel M^*)$</p> <p>04 if $c^* \neq H'(W^* \parallel M^*)$ then abort</p> <p>05 if $\llbracket M^* \notin \mathcal{M} \rrbracket \wedge \forall(pk, W^*, c^*, Z^*)$ then return (M^*, σ^*)</p> <p>06 abort</p>

Figure 9: Adversary B against UF-NMA security of SIG with quantum access to random oracle H' . The oracles SIGN_1 and H simulated by B are defined as in game G_2 of Figure 8.

We distinguish two cases. If $(M^*, \cdot) \notin \mathcal{M}$ then we are in the situation that the adversary did not query a signature on M^* and we can use the same argument as in standard unforgeability to argue $|\Pr[G_2^A \Rightarrow 1] - \Pr[G_1^A \Rightarrow 1]| \leq 2^{-\alpha+1}$. It leaves to handle the case $(M^*, \cdot) \in \mathcal{M}$, i.e., the adversary obtained a signatures $\sigma_{M^*} = (W_{M^*}, Z_{M^*})$ on message M^* and submits a correct forgery $\sigma^* = (W^*, Z^*)$ satisfying $W^* = W_{M^*}$ and $Z^* \neq Z_{M^*}$. The problem of finding values (W^*, c^*, Z_{M^*}, Z^*) with two accepting transcripts (W^*, c^*, Z^*) and (W^*, c^*, Z_{M^*}) is exactly bounded by the advantage of an adversary C against the CUR experiment, i.e., $|\Pr[G_2^A \Rightarrow 1] - \Pr[G_1^A \Rightarrow 1]| \leq \text{Adv}_{\text{ID}}^{\text{CUR}}(\text{C})$.

In combination this proves

$$|\Pr[G_2^A \Rightarrow 1] - \Pr[G_1^A \Rightarrow 1]| \leq 2^{-\alpha+1} + \text{Adv}_{\text{ID}}^{\text{CUR}}(\text{C}).$$

Finally, a straightforward modification of adversary B against UF-NMA security to account for the strong unforgeability check proves

$$\Pr[G_2^A \Rightarrow 1] = \text{Adv}_{\text{SIG}}^{\text{UF-NMA}}(\text{B})$$

and completes proof of sUF-CMA₁ security.

The running times $\text{Time}(\text{B})$ and $\text{Time}(\text{C})$ can be derived as above. \square

The following theorem shows that we can also prove directly UF-CMA security of SIG, but (in terms of the running time) the reduction is less tight than the one of Theorem 3.2.

Theorem 3.3 *Assume the identification scheme ID is ε_{zk} -perfect naHVZK and has α bits of min entropy. For any UF-CMA (sUF-CMA) quantum adversary A that issues at most Q_H queries to the quantum random oracle $|H\rangle$ and Q_S classical queries to the signing oracle SIGN, there exists a quantum adversary B against UF-NMA security making Q_H queries to its own quantum random oracle (and a quantum adversary C against CUR) such that*

$$\begin{aligned} \text{Adv}_{\text{SIG}}^{\text{UF-CMA}}(\text{A}) &\leq \text{Adv}_{\text{SIG}}^{\text{UF-NMA}}(\text{B}) + Q_S \cdot 2^{-\alpha+1} + \kappa_m Q_S \cdot \varepsilon_{\text{zk}}, \\ \text{Adv}_{\text{SIG}}^{\text{sUF-CMA}}(\text{A}) &\leq \text{Adv}_{\text{SIG}}^{\text{UF-NMA}}(\text{B}) + Q_S \cdot 2^{-\alpha+1} + \kappa_m Q_S \cdot \varepsilon_{\text{zk}} + \text{Adv}_{\text{ID}}^{\text{CUR}}(\text{C}), \end{aligned}$$

and $\text{Time}(\text{B}) = \text{Time}(\text{C}) = \text{Time}(\text{A}) + \kappa_m Q_H Q_S$.

The proof of Theorem 3.3 is similar to the one of Theorem 3.2 and postponed to Appendix A.2.

3.2.2 From Lossiness to UF-NMA

Theorem 3.4 *Assume the identification scheme is lossy and ε_{ls} -lossy sound. For any UF-NMA quantum adversary A that issues at most Q_H queries to the quantum random oracle $|H\rangle$, there exists a quantum adversary B against LOSS such that*

$$\text{Adv}_{\text{SIG}}^{\text{UF-NMA}}(\text{A}) \leq \text{Adv}_{\text{ID}}^{\text{LOSS}}(\text{B}) + 8(Q_H + 1)^2 \cdot \varepsilon_{\text{ls}},$$

and $\text{Time}(\text{B}) = \text{Time}(\text{A}) + Q_H \approx \text{Time}(\text{A})$.

Proof. Let A be an adversary against the UF-NMA security of SIG, issuing at most Q_H quantum queries to $|H\rangle$. Consider the games given in Figure 10.

GAME G_0-G_1	
01 $(pk, sk) \leftarrow \text{IGen}(\text{par})$	// G_0
02 $pk \leftarrow \text{LossyGen}(\text{par})$	// G_1
03 $(M^*, \sigma^*) \leftarrow \text{A}^{\text{H}}(pk)$	
04 Parse $\sigma^* = (W^*, Z^*)$	
05 $c^* := \text{H}(W^* \parallel M^*)$	
06 return $\text{V}(pk, W^*, c^*, Z^*)$	

Figure 10: Games G_0-G_1 for the proof of Theorem 3.4.

GAME G_0 . Since game G_0 is the original UF-NMA game,

$$\Pr[G_0^{\text{A}} \Rightarrow 1] = \text{Adv}_{\text{SIG}}^{\text{UF-NMA}}(\text{A}) .$$

GAME G_1 . In this game, the public key pk is changed to lossy mode. Clearly, there exists an adversary B simulating H by a $2Q_{\text{H}}$ -wise independent hash function such that

$$|\Pr[G_1^{\text{A}} \Rightarrow 1] - \Pr[G_0^{\text{A}} \Rightarrow 1]| \leq \text{Adv}_{\text{ID}}^{\text{LOSS}}(\text{B}) .$$

Adversary C_1	Adversary $\text{C}_2^{(g)}$
01 $pk \leftarrow \text{LossyGen}(\text{par})$	08 $(M^*, \sigma^*) \leftarrow \text{A}^{\text{H}}(pk)$
02 Pick $2Q_{\text{H}}$ -wise independent $f_{2Q_{\text{H}}}$	09 Parse $\sigma^* = (W^*, Z^*)$
03 for each $W \in \text{WSet}$ do	10 $c^* := \text{H}(W^* \parallel M^*)$
04 compute set $\text{ChGOOD}_{pk}(W) \subseteq \text{ChSet}$	11 if $\text{V}(pk, W^*, c^*, Z^*) = 1$ return $(W^* \parallel M^*)$
05 $\lambda_{pk}(W) := \text{ChGOOD}_{pk}(W) / \text{ChSet} $	12 else return \perp
06 for each $M \in \text{MSet}$ set $\lambda_{pk}(W \parallel M) := \lambda_{pk}(W)$	
07 return $(\lambda_{pk}(W \parallel M))_{W \in \text{WSet}, M \in \text{MSet}}$	
$\text{H}(W \parallel M)$	// quantum access
13 $y := g(W \parallel M)$	
14 if $y = 1$: $c := \text{Uni}(\text{ChGOOD}_{pk}(W); f_{2Q_{\text{H}}}(W \parallel M))$	
15 if $y = 0$: $c := \text{Uni}(\text{ChSet} \setminus \text{ChGOOD}_{pk}(W); f_{2Q_{\text{H}}}(W \parallel M))$	
16 return c	

Figure 11: Adversary $\text{C} = (\text{C}_1, \text{C}_2)$ in game GSPB for the proof of Theorem 3.4. The set of good challenges $\text{ChGOOD}_{pk}(W)$ is defined in Equation (6).

Finally, we will reduce a successful A in game G_1 to the generic search problem GSPB to show

$$\Pr[G_1^{\text{A}} \Rightarrow 1] \leq 8(Q_{\text{H}} + 1)^2 \varepsilon_{\text{Is}}. \quad (5)$$

For a finite set S , let $\text{Uni}(S)$ be a probabilistic algorithm that returns uniform $x \leftarrow S$ and recall that $x := \text{Uni}(S; r)$ denotes the deterministic execution of $\text{Uni}(S)$ using explicitly given random tape r . To prove Equation (5), consider the unbounded adversary $\text{C} = (\text{C}_1, \text{C}_2)$ defined in Figure 11 that is executed in the generic search game GSPB, making at most Q_{H} quantum queries to the oracle $|g(\cdot)\rangle$. First note that computing the probabilities $\lambda_{pk}(W \parallel M) = \lambda_{pk}(W)$ in line 05 for all $W \in \text{WSet}$ and $M \in \text{MSet}$ may take exponential time but since C is computationally unbounded it does not matter.

To analyze C 's success probability in game GSPB, we first fix a public-key pk . Now consider some $W \parallel M$ with non-zero amplitude as part of a query to quantum random oracle H . Set $\text{ChGOOD}_{pk}(W)$ of “good challenges” is defined as

$$\text{ChGOOD}_{pk}(W) := \{c \in \text{ChSet} \mid \exists Z \in \text{ZSet} : \text{V}(pk, W, c, Z) = 1\}. \quad (6)$$

That is, the set $\text{ChGOOD}_{pk}(W)$ contains all challenges c for which there exists a possible response Z to make (W, c, Z) a valid transcript (with respect to pk). By definition of GSPB, each query to oracle $g(W \parallel M)$ returns $y = 1$ with probability $\lambda_{pk}(W \parallel M) = |\text{ChGOOD}_{pk}(W)|/|\text{ChSet}|$. Hence, the output distribution of $\text{H}(W \parallel M)$ sampled in lines 14 and 15 is uniform over ChSet , as in game G_1 . Consistency

of \mathbf{H} is assured by deriving the randomness to sample c in case $y = 0$ (lines 14 and 15) using fixed random coins $f_{2Q_H}(W \parallel M)$, derived by a $2Q_H$ -wise independent hash function f_{2Q_H} (which looks like a perfectly random function to \mathbf{A}).

Now consider \mathbf{A} 's forgery $\sigma^* = (W^*, Z^*)$ on message M^* and define $c^* := \mathbf{H}(W^* \parallel M^*)$. If the signature is valid (i.e., $\mathbf{V}(pk, W^*, c^*, Z^*) = 1$), then clearly c^* is a good challenge from set $\text{ChGOOD}_{pk}(W^*)$ which implies $g(W^* \parallel M^*) = 1$. This proves

$$\Pr[G_1 \Rightarrow 1 \mid pk] = \Pr[\text{GSPB}_{\lambda_{pk}}^c \Rightarrow 1 \mid pk] \leq 8(Q_H + 1)^2 \lambda_{pk}, \quad (7)$$

where

$$\lambda_{pk} = \max_{W \in \text{WSet}, M \in \text{MSet}} \lambda_{pk}(W \parallel M)$$

Averaging Equation (7) over $pk \leftarrow \text{LossyGen}$ we finally obtain

$$\Pr[G_1 \Rightarrow 1] \leq 8(Q_H + 1)^2 \cdot \mathbf{E}_{pk}[\lambda_{pk}] \leq 8(Q_H + 1)^2 \varepsilon_{\text{ls}},$$

where the last inequality uses Equation (3) for the optimal adversary. \square

4 Dilithium-QROM

In this section, we present a modification of the Dilithium digital signature scheme [DLL⁺17] whose security is based on MLWE in the QROM. We also present a new security proof of the original Dilithium that shows it to be tightly-secure in the QROM based on a different non-interactive assumption. Since Dilithium is a highly-optimized version of a scheme constructed via the ‘‘Fiat-Shamir with Aborts’’ framework [Lyu09], its details may be somewhat overwhelming to readers who are not already comfortable with such constructions. For this reason, we present a much simpler version of the signature scheme without any optimizations in Appendix B. The reader who wishes to get a feel for how the techniques of the previous section apply to lattice-based schemes is thus encouraged to read this section after familiarizing oneself with the notation in Section 4.1.1.

4.1 Preliminaries

4.1.1 Rings and Distributions

We let R and R_q respectively denote the rings $\mathbb{Z}[X]/(X^n + 1)$ and $\mathbb{Z}_q[X]/(X^n + 1)$, for an integer q . We will assume that $q \equiv 5 \pmod{8}$, as such a choice of q ensures that all polynomials in R_q with coefficients less than $\sqrt{q/2}$ have an inverse in the ring [LN17, Lemma 2.2]. This property is crucial to our security proof. Regular font letters denote elements in R or R_q (which includes elements in \mathbb{Z} and \mathbb{Z}_q) and bold lower-case letters represent column vectors with coefficients in R or R_q . By default, all vectors will be column vectors. Bold upper-case letters are matrices.

MODULAR REDUCTIONS. For an even (resp. odd) positive integer α , we define $r' = r \bmod^\pm \alpha$ to be the unique element r' in the range $-\frac{\alpha}{2} < r' \leq \frac{\alpha}{2}$ (resp. $-\frac{\alpha-1}{2} \leq r' \leq \frac{\alpha-1}{2}$) such that $r' = r \bmod \alpha$. We will sometimes refer to this as a *centered* reduction modulo q . For any positive integer α , we define $r' = r \bmod^+ \alpha$ to be the unique element r' in the range $0 \leq r' < \alpha$ such that $r' = r \bmod \alpha$. When the exact representation is not important, we simply write $r \bmod \alpha$.

SIZES OF ELEMENTS. For an element $w \in \mathbb{Z}_q$, we write $\|w\|_\infty$ to mean $|w \bmod^\pm q|$. We now define the ℓ_∞ and ℓ_2 norms for $w = w_0 + w_1X + \dots + w_{n-1}X^{n-1} \in R$:

$$\|w\|_\infty = \max_i \|w_i\|_\infty, \quad \|w\| = \sqrt{\|w_0\|_\infty^2 + \dots + \|w_{n-1}\|_\infty^2}.$$

Similarly, for $\mathbf{w} = (w_1, \dots, w_k) \in R^k$, we define

$$\|\mathbf{w}\|_\infty = \max_i \|w_i\|_\infty, \quad \|\mathbf{w}\| = \sqrt{\|w_1\|^2 + \dots + \|w_k\|^2}.$$

We will write S_η to denote all elements $w \in R$ such that $\|w\|_\infty \leq \eta$.

EXTENDABLE OUTPUT FUNCTION. Suppose that Sam is an extendable output function, that is a function on bit strings in which the output can be extended to any desired length. If we would like Sam to take as input x and then produce a value y that is distributed according to distribution S (or uniformly over a set S), we write $y \sim S := \text{Sam}(x)$. It is important to note that this procedure is completely deterministic: a given x will always produce the same y . For simplicity we assume that the output distribution of Sam is perfect, whereas in practice Sam will be implemented using random oracles and produce an output that is statistically close to the perfect distribution. If K is a secret key, then $\text{Sam}(K||x)$ is a pseudo-random function from $\{0, 1\}^* \rightarrow \{0, 1\}^*$.

THE CHALLENGE SPACE. The challenge space in our identification and signature schemes needs to be a subset of the ring R , have size a little larger than 2^{256} , and consist of polynomials with small norms. In this paper, the dimension n of the ring R will be taken to be 512,³ and so we will define the challenge space accordingly as

$$\text{ChSet} := \{c \in R \mid \|c\|_\infty = 1 \text{ and } \|c\| = \sqrt{46}\}. \quad (8)$$

In other words, ChSet consists of elements in R with $-1/0/1$ coefficients that have exactly 46 non-zero coefficients. The size of this set is $\binom{n}{46} \cdot 2^{46}$, which for $n = 512$ is greater than 2^{265} .

THE MLWE ASSUMPTION. For integers m, k , and a probability distribution $D : R_q \rightarrow [0, 1]$, we say that the advantage of algorithm A in solving the decisional $\text{MLWE}_{m,k,D}$ problem over the ring R_q is

$$\begin{aligned} \text{Adv}_{m,k,D}^{\text{MLWE}} := & \left| \Pr[A(\mathbf{A}, \mathbf{t}) \Rightarrow 1 \mid \mathbf{A} \leftarrow R_q^{m \times k}; \mathbf{t} \leftarrow R_q^m] \right. \\ & \left. - \Pr[A(\mathbf{A}, \mathbf{A}\mathbf{s}_1 + \mathbf{s}_2) \Rightarrow 1 \mid \mathbf{A} \leftarrow R_q^{m \times k}; \mathbf{s}_1 \leftarrow D^k; \mathbf{s}_2 \leftarrow D^m] \right|. \end{aligned}$$

The MLWE assumption states that the above advantage is negligible for all polynomial-time algorithms A . This assumption was introduced in [LS15], and is generalization of the LWE assumption from [Reg05]. The Ring-LWE assumption [LPR10] is a special case of MLWE where $k = 1$. Analogously to LWE and Ring-LWE, it was shown in [LS15] that solving the MLWE problem for certain parameters is as hard as solving certain worst-case problems in certain algebraic lattices.

4.1.2 Summary of Supporting Algorithms

To reduce the size of the public key, we will need some simple algorithms that extract “higher-order” and “lower-order” bits of elements in \mathbb{Z}_q . The goal is that when given an arbitrary element $r \in \mathbb{Z}_q$ and another small element $z \in \mathbb{Z}_q$, we would like to be able to recover the higher order bits of $r + z$ without needing to store z . We therefore define algorithms that take r, z and produce a 1-bit hint h that allows one to compute the higher order bits of $r + z$ just using r and h . This hint is essentially the “carry” caused by z in the addition. The algorithms are exactly as in [DLL⁺17], and we repeat them for convenience in Figure 12. The algorithms are described as working on integers modulo q , but are extended to polynomials in R_q by simply being applied individually to each coefficient.

The below Lemmas recall the crucial properties of these supporting algorithms that are necessary for the correctness and security of our scheme.

Lemma 4.1 *Suppose that q and α are positive integers satisfying $q > 2\alpha$, $q \equiv 1 \pmod{\alpha}$ and α even. Let \mathbf{r} and \mathbf{z} be vectors of elements in R_q where $\|\mathbf{z}\|_\infty \leq \alpha/2$, and let \mathbf{h}, \mathbf{h}' be vectors of bits. Then the HighBits_q , MakeHint_q , and UseHint_q algorithms satisfy the following properties:*

1. $\text{UseHint}_q(\text{MakeHint}_q(\mathbf{z}, \mathbf{r}, \alpha), \mathbf{r}, \alpha) = \text{HighBits}_q(\mathbf{r} + \mathbf{z}, \alpha)$.
2. Let $\mathbf{v}_1 = \text{UseHint}_q(\mathbf{h}, \mathbf{r}, \alpha)$. Then $\|\mathbf{r} - \mathbf{v}_1 \cdot \alpha\|_\infty \leq \alpha + 1$.
3. For any \mathbf{h}, \mathbf{h}' , if $\text{UseHint}_q(\mathbf{h}, \mathbf{r}, \alpha) = \text{UseHint}_q(\mathbf{h}', \mathbf{r}, \alpha)$, then $\mathbf{h} = \mathbf{h}'$.

Lemma 4.2 *If $\|\mathbf{s}\|_\infty \leq \beta$ and $\|\text{LowBits}_q(\mathbf{r}, \alpha)\|_\infty < \alpha/2 - \beta$, then*

$$\text{HighBits}_q(\mathbf{r}, \alpha) = \text{HighBits}_q(\mathbf{r} + \mathbf{s}, \alpha).$$

Power2Round_q(r, d)	Decompose_q(r, α)
01 $r := r \bmod^+ q$	12 $r := r \bmod^+ q$
02 $r_0 := r \bmod^{\pm} 2^d$	13 $r_0 := r \bmod^{\pm} \alpha$
03 return $(r - r_0)/2^d$	14 if $r - r_0 = q - 1$
UseHint_q(h, r, α)	15 then $r_1 := 0; r_0 := r_0 - 1$
04 $m := (q - 1)/\alpha$	16 else $r_1 := (r - r_0)/\alpha$
05 $(r_1, r_0) := \text{Decompose}_q(r, \alpha)$	17 return (r_1, r_0)
06 if $h = 1$ and $r_0 > 0$ return $(r_1 + 1) \bmod^+ m$	HighBits_q(r, α)
07 if $h = 1$ and $r_0 \leq 0$ return $(r_1 - 1) \bmod^+ m$	18 $(r_1, r_0) := \text{Decompose}_q(r, \alpha)$
08 return r_1	19 return r_1
MakeHint_q(z, r, α)	LowBits_q(r, α)
09 $r_1 := \text{HighBits}_q(r, \alpha)$	20 $(r_1, r_0) := \text{Decompose}_q(r, \alpha)$
10 $v_1 := \text{HighBits}_q(r + z, \alpha)$	21 return r_0
11 return $\llbracket r_1 \neq v_1 \rrbracket$	

Figure 12: Supporting algorithms for Dilithium and Dilithium-QROM.

4.2 The Identification Protocol

The constituting algorithms of our identification protocol $\text{ID} = (\text{IGen}, \text{P}_1, \text{P}_2, \text{V})$ are described in Figure 13 with the concrete parameters $\text{par} = (q, n, k, \ell, d, \gamma, \gamma', \eta, \beta)$ given later in Table 1.

KEY GENERATION. The key generation proceeds by choosing a random 256-bit seed ρ and expanding into a matrix $\mathbf{A} \in R_q^{k \times \ell}$ by an extendable output function Sam modeled as a random oracle. The secret keys $(\mathbf{s}_1, \mathbf{s}_2) \in S_\eta^\ell \times S_\eta^k$ have uniformly random coefficients between $-\eta$ and η (inclusively). The value $\mathbf{t} = \mathbf{A}\mathbf{s}_1 + \mathbf{s}_2$ is then computed. The public key that is needed for verification is (ρ, \mathbf{t}_1) with \mathbf{t}_1 output by the $\text{Power2Round}_q(\mathbf{t}, d)$ algorithm in Figure 12 (we have $\mathbf{t} = \mathbf{t}_1 \cdot 2^d + \mathbf{t}_0$ for some small \mathbf{t}_0), while the secret key is $(\rho, \mathbf{s}_1, \mathbf{s}_2, \mathbf{t}_0)$.

While the verifier never needs the value \mathbf{t}_0 (and thus it does not need to be included in the public key of the actual scheme), we do need this value in order to simulate transcripts (see Section 4.3.1). Thus the security of our scheme is based on the fact that the adversary gets \mathbf{t}_1 and \mathbf{t}_0 , whereas in reality he only gets \mathbf{t}_1 .

The set ChSet is defined as in Equation (8), and $\text{ZSet} = S_{\gamma' - \beta - 1}^\ell \times \{0, 1\}^k$. The set of commitments WSet is defined as $\text{WSet} = \{\mathbf{w}_1 : \exists \mathbf{y} \in S_{\gamma' - 1}^\ell \text{ s.t. } \mathbf{w}_1 = \text{HighBits}_q(\mathbf{A}\mathbf{y}, 2\gamma)\}$.

PROTOCOL EXECUTION. The prover starts the identification protocol by reconstructing \mathbf{A} from the random seed ρ . The next step has the prover sample $\mathbf{y} \leftarrow S_{\gamma' - 1}^\ell$ and then compute $\mathbf{w} = \mathbf{A}\mathbf{y}$. He then writes $\mathbf{w} = 2\gamma \cdot \mathbf{w}_1 + \mathbf{w}_0$, with \mathbf{w}_0 between $-\gamma$ and γ (inclusively), and then sends \mathbf{w}_1 to the verifier. The verifier generates a random challenge $c \leftarrow \text{ChSet}$ and sends it to the prover. The prover computes $\mathbf{z} = \mathbf{y} + c\mathbf{s}$. If $\mathbf{z} \notin S_{\gamma' - \beta - 1}^\ell$, then the prover sets his response to \perp . He also replies with \perp if $\text{LowBits}_q(\mathbf{w} - c\mathbf{s}_2, 2\gamma) \notin S_{\gamma - \beta - 1}^k$. This part of the protocol is necessary for security – it makes sure that \mathbf{z} does not leak anything about the secret key $\mathbf{s}_1, \mathbf{s}_2$.

If the checks pass and a \perp is not sent, then it can be shown (see Section 4.3.2) that $\text{HighBits}_q(\mathbf{A}\mathbf{z} - c\mathbf{t}, 2\gamma) = \mathbf{w}_1$. At this point, if the verifier knew the entire element \mathbf{t} and (\mathbf{z}, c) , he could have recovered \mathbf{w}_1 and checked that $\|\mathbf{z}\|_\infty < \gamma' - \beta$ and that the high-order bits of $\mathbf{A}\mathbf{z} - c\mathbf{t}$ are indeed \mathbf{w}_1 . However, since we want to compress the size of the public key, the verifier only knows \mathbf{t}_1 . Hence, the signer needs to provide a “hint” \mathbf{h} which will allow the verifier to compute $\text{HighBits}_q(\mathbf{A}\mathbf{z} - c\mathbf{t}, 2\gamma)$.

The verifier checks whether $\|\mathbf{z}\|_\infty < \gamma' - \beta$ and that $\mathbf{A}\mathbf{z} - c\mathbf{t}_1 \cdot 2^d$ together with the hint \mathbf{h} allow him to reconstruct \mathbf{w}_1 . We should point out that in the identification scheme it is actually not necessary for the verifier to be able to recover exactly \mathbf{w}_1 . He could have simply checked that $\mathbf{A}\mathbf{z} - c\mathbf{t}_1 \cdot 2^d \approx \mathbf{w}_1$ and this would be good enough for security. The reason that we want the verifier to be able to exactly recover \mathbf{w}_1 is to make the ID scheme *commitment-recoverable* and be able to reduce the communication size in the Fiat-Shamir transform (see Section 3.1).

³In Section 4.5, we will also discuss a scheme where $n = 256$. For that scheme the challenge space consists of 60 ± 1 's.

<p>IGen(par)</p> <pre> 01 $\rho \leftarrow \{0, 1\}^{256}$ 02 $\mathbf{A} \leftarrow R_q^{k \times \ell} := \text{Sam}(\rho)$ 03 $(\mathbf{s}_1, \mathbf{s}_2) \leftarrow S_\eta^\ell \times S_\eta^k$ 04 $\mathbf{t} := \mathbf{A}\mathbf{s}_1 + \mathbf{s}_2$ 05 $\mathbf{t}_1 := \text{Power2Round}_q(\mathbf{t}, d)$ 06 $\mathbf{t}_0 := \mathbf{t} - \mathbf{t}_1 \cdot 2^d$ 07 $pk = (\rho, \mathbf{t}_1, \mathbf{t}_0)$ 08 $sk = (\rho, \mathbf{s}_1, \mathbf{s}_2, \mathbf{t}_0)$ 09 return (pk, sk) </pre> <p>V$(pk, W = \mathbf{w}_1, c, Z = (\mathbf{z}, \mathbf{h}))$</p> <pre> 20 return $[\ \mathbf{z}\ _\infty < \gamma' - \beta]$ and $[\mathbf{w}_1 = \text{UseHint}_q(\mathbf{h}, \mathbf{A}\mathbf{z} - \mathbf{c}\mathbf{t}_1 \cdot 2^d, 2\gamma)]$ </pre>	<p>P₁(sk)</p> <pre> 10 $\mathbf{A} \leftarrow R_q^{k \times \ell} := \text{Sam}(\rho)$ 11 $\mathbf{y} \leftarrow S_{\gamma'-1}^\ell$ 12 $\mathbf{w} := \mathbf{A}\mathbf{y}$ 13 $\mathbf{w}_1 := \text{HighBits}_q(\mathbf{w}, 2\gamma)$ 14 return $(W = \mathbf{w}_1, St = (\mathbf{w}, \mathbf{y}))$ </pre> <p>P₂(sk, W = w₁, c, St = (w, y))</p> <pre> 15 $\mathbf{z} := \mathbf{y} + \mathbf{c}\mathbf{s}_1$ 16 if $\ \mathbf{z}\ _\infty \geq \gamma' - \beta$ or $\ \text{LowBits}_q(\mathbf{w} - \mathbf{c}\mathbf{s}_2, 2\gamma)\ _\infty \geq \gamma - \beta$ 17 then $(\mathbf{z}, \mathbf{h}) := \perp$ 18 else $\mathbf{h} := \text{MakeHint}_q(-\mathbf{c}\mathbf{t}_0, \mathbf{w} - \mathbf{c}\mathbf{s}_2 + \mathbf{c}\mathbf{t}_0, 2\gamma)$ 19 return $Z = (\mathbf{z}, \mathbf{h})$ </pre>
--	--

Figure 13: Our ID scheme – a concrete instantiation based on the hardness of the MLWE problem of the commitment-recoverable (Definition 2.4) canonical identification scheme in Figure 4. The \mathbf{t}_0 part of the public key is assumed to be known by the adversary in the security proofs, but is not needed by the verifier for verification. Thus in the real scheme, \mathbf{t}_0 would not be included as part of the public key.

4.3 Security Properties

4.3.1 Non Abort Honest Verifier Zero-Knowledge

In this section, we will show that ID is perfectly naHVZK, i.e., the distribution of the output of the Trans algorithm (Figure 14, left) that uses the secret key as input is exactly that of the Sim algorithm (Figure 14, right) that uses only the public key as input.

<p>Algorithm Trans(sk):</p> <pre> 01 $\mathbf{A} \leftarrow R_q^{k \times \ell} := \text{Sam}(\rho)$ 02 $\mathbf{y} \leftarrow S_{\gamma'-1}^\ell$ 03 $\mathbf{w} := \mathbf{A}\mathbf{y}$ 04 $\mathbf{w}_1 := \text{HighBits}_q(\mathbf{w}, 2\gamma)$ 05 $c \leftarrow \text{ChSet}$ 06 $\mathbf{z} := \mathbf{y} + \mathbf{c}\mathbf{s}_1$ 07 if $\ \mathbf{z}\ _\infty \geq \gamma' - \beta$ then return \perp 08 if $\ \text{LowBits}_q(\mathbf{w} - \mathbf{c}\mathbf{s}_2, 2\gamma)\ _\infty \geq \gamma - \beta$ then return \perp 09 $\mathbf{h} := \text{MakeHint}_q(-\mathbf{c}\mathbf{t}_0, \mathbf{w} - \mathbf{c}\mathbf{s}_2 + \mathbf{c}\mathbf{t}_0, 2\gamma)$ 10 return $(c, (\mathbf{z}, \mathbf{h}))$ </pre>	<p>Algorithm Sim(pk):</p> <pre> 11 $\mathbf{A} \leftarrow R_q^{k \times \ell} := \text{Sam}(\rho)$ 12 with probability $1 - \frac{ S_{\gamma'-\beta-1}^\ell }{ S_{\gamma'-1}^\ell }$, return \perp 13 $\mathbf{z} \leftarrow S_{\gamma'-\beta-1}^\ell$ 14 $c \leftarrow \text{ChSet}$ 15 if $\ \text{LowBits}_q(\mathbf{A}\mathbf{z} - \mathbf{c}\mathbf{t}, 2\gamma)\ _\infty \geq \gamma - \beta$ 16 then return \perp 17 $\mathbf{h} := \text{MakeHint}_q(-\mathbf{c}\mathbf{t}_0, \mathbf{A}\mathbf{z} - \mathbf{c}\mathbf{t} + \mathbf{c}\mathbf{t}_0, 2\gamma)$ 18 return $(c, (\mathbf{z}, \mathbf{h}))$ </pre>
---	---

Figure 14: Left: a real transcript output by the transcript algorithm Trans(sk); Right: a simulated transcript output by the Sim(pk) algorithm.

Lemma 4.3 *If $\beta \geq \max_{s \in S_\eta, c \in \text{ChSet}} \|cs\|_\infty$, then ID is perfectly naHVZK.*

Proof. Let $(\mathbf{s}_1, \mathbf{s}_2) \in S_\eta^\ell \times S_\eta^k$ be any polynomials satisfying $\mathbf{A}\mathbf{s}_1 + \mathbf{s}_2 = \mathbf{t}$. We will show that the output distributions of Trans and Sim from Figure 14 are identical.

For any $\mathbf{z} \in S_{\gamma'-\beta-1}^k$, let us compute the probability of it being generated in line 06 of Trans. For any $c \in \text{ChSet}$, we have

$$\Pr_{\mathbf{y} \leftarrow S_{\gamma'-1}^\ell} [\mathbf{y} + \mathbf{c}\mathbf{s}_1 = \mathbf{z}] = \Pr_{\mathbf{y} \leftarrow S_{\gamma'-1}^\ell} [\mathbf{y} = \mathbf{z} - \mathbf{c}\mathbf{s}_1].$$

Notice that because $\|\mathbf{c}\mathbf{s}_1\|_\infty \leq \beta$, we know that $\mathbf{z} - \mathbf{c}\mathbf{s}_1 \in S_{\gamma'-1}^\ell$. Thus

$$\Pr_{\mathbf{y} \leftarrow S_{\gamma'-1}^\ell} [\mathbf{y} = \mathbf{z} - \mathbf{c}\mathbf{s}_1] = 1/|S_{\gamma'-1}^\ell|. \quad (9)$$

Therefore every $\mathbf{z} \in S_{\gamma'-\beta-1}^\ell$ has an equal probability of being generated. Furthermore, the probability of producing a $\mathbf{z} \in S_{\gamma'-\beta-1}^\ell$, and thus not returning \perp in Line 07 of Trans, is exactly $\frac{|S_{\gamma'-\beta-1}^\ell|}{|S_{\gamma'-1}^\ell|}$. Thus after the completion of Line 07, either \perp has been returned (with probability $1 - \frac{|S_{\gamma'-\beta-1}^\ell|}{|S_{\gamma'-1}^\ell|}$), or the distribution of (c, \mathbf{z}) is uniform in $\text{ChSet} \times S_{\gamma'-\beta-1}^\ell$. This is exactly the same distribution as after Line 14 of Sim.

To complete the proof, we note that

$$\mathbf{w} - c\mathbf{s}_2 = \mathbf{A}\mathbf{y} - c\mathbf{s}_2 = \mathbf{A}(\mathbf{z} - c\mathbf{s}_1) - c\mathbf{s}_2 = \mathbf{A}\mathbf{z} - c\mathbf{t}, \quad (10)$$

and therefore all the steps in Trans after Line 07 are identical to those after Line 14 of Sim. \square

4.3.2 Correctness

In this section, we compute the probability that the Prover does not send \perp and then show that the verification procedure will always accept a transcript when the Prover does not send \perp .

Lemma 4.4 *If $\beta \geq \max_{s \in S_\eta, c \in \text{ChSet}} \|cs\|_\infty$ then ID has correctness error $\delta \approx 1 - \exp(-\beta n \cdot (k/\gamma + \ell/\gamma'))$.*

Proof. Let $\mathbf{s}_1, \mathbf{s}_2 \in S_\eta^k$ be any polynomials satisfying $\mathbf{A}\mathbf{s}_1 + \mathbf{s}_2 = \mathbf{t}$. We first show

$$\Pr_{\mathbf{y} \leftarrow S_{\gamma'-1}^\ell, c \leftarrow \text{ChSet}} [(\mathbf{z}, \mathbf{h}) \neq \perp] \approx \exp(-\beta n \cdot (k/\gamma + \ell/\gamma')). \quad (11)$$

Since we showed in Lemma 4.3 that Sim outputs \perp with the same probability as Trans (and therefore with the same probability as in the identification scheme), we can simply compute the probability that Sim does not output \perp .

The probability that \perp is not output in Line 12 of Sim (Figure 14) is

$$\frac{|S_{\gamma'-\beta-1}^\ell|}{|S_{\gamma'-1}^\ell|} = \left(\frac{2(\gamma' - \beta) - 1}{2\gamma' - 1} \right)^{n\ell} > \left(1 - \frac{\beta}{\gamma'} \right)^{n\ell} \approx e^{-\beta n \ell / \gamma'}, \quad (12)$$

where the approximate equality is due to the fact that $\beta \ll \gamma$.

If we then heuristically assume that for a uniformly-random $\mathbf{z} \in S_{\gamma'-\beta-1}^k$, the distribution of $\mathbf{A}\mathbf{z} - c\mathbf{t} \bmod 2\gamma$ is approximately uniform, then we also have that the probability that \perp is not output in Line 16 is

$$\Pr_{\mathbf{z} \leftarrow S_{\gamma'-\beta-1}^k} [\|\text{LowBits}_q(\mathbf{A}\mathbf{z} - c\mathbf{t}, 2\gamma)\|_\infty < \gamma - \beta] \approx \frac{|S_{\gamma-\beta-1}^k|}{|S_{\gamma-1}^k|} \approx e^{-\beta n k / \gamma}. \quad (13)$$

Multiplying the bounds in Equation (12) and Equation (13) proves Equation (11).

To complete the proof of the lemma, it remains to show that if $(\mathbf{z}, \mathbf{h}) \neq \perp$, then the verification procedure will always accept. Assume $(\mathbf{z}, \mathbf{h}) \neq \perp$. It's clear that the verifier's check $\|\mathbf{z}\|_\infty < \gamma' - \beta$ will always pass. We will now show that the second check will pass as well. Because

$$\mathbf{w} - c\mathbf{s}_2 = \mathbf{A}\mathbf{y} - c\mathbf{s}_2 = \mathbf{A}(\mathbf{z} - c\mathbf{s}_1) - c\mathbf{s}_2 = \mathbf{A}\mathbf{z} - c\mathbf{t} = \mathbf{A}\mathbf{z} - c\mathbf{t}_0 - c\mathbf{t}_1 \cdot 2^d,$$

we can rewrite the MakeHint call of the prover as

$$\mathbf{h} := \text{MakeHint}_q(-c\mathbf{t}_0, \mathbf{w} - c\mathbf{s}_2 + c\mathbf{t}_0, 2\gamma) = \text{MakeHint}_q(-c\mathbf{t}_0, \mathbf{A}\mathbf{z} - c\mathbf{t}_1 \cdot 2^d, 2\gamma).$$

Since $\|c\mathbf{t}_0\|_\infty < \gamma$, by Lemma 4.1 we know that the verifier computes

$$\begin{aligned} \text{UseHint}_q(\mathbf{h}, \mathbf{A}\mathbf{z} - c\mathbf{t}_1 \cdot 2^d, 2\gamma) &= \text{HighBits}_q(\mathbf{A}\mathbf{z} - c\mathbf{t}_1 \cdot 2^d, 2\gamma) \\ &= \text{HighBits}_q(\mathbf{w} - c\mathbf{s}_2, 2\gamma). \end{aligned} \quad (14)$$

Since $\|\text{LowBits}_q(\mathbf{w} - c\mathbf{s}_2, 2\gamma)\|_\infty < \gamma - \beta$, we know from Lemma 4.2 that

$$\text{HighBits}_q(\mathbf{w} - c\mathbf{s}_2, 2\gamma) = \text{HighBits}_q(\mathbf{w}, 2\gamma) = \mathbf{w}_1.$$

Therefore the verifier will correctly compute $\mathbf{w}_1 := \text{UseHint}_q(\mathbf{h}, \mathbf{A}\mathbf{z} - c\mathbf{t}_1 \cdot 2^d, 2\gamma)$ and will accept a valid prover. \square

4.3.3 Lossyness

In this section, we analyze the scheme in which the public key is generated uniformly at random, as in algorithm `LossyGen` of Figure 15, rather than as in `IGen` of Figure 13. Our goal is to show that even if the prover is computationally unbounded, he only has approximately a $1/|\text{ChSet}|$ probability of making the verifier accept during each run of the identification scheme. This will show that the probability in Equation (3) is upper-bounded by approximately $1/|\text{ChSet}|$.

<pre> LossyGen(par) 01 $\rho \leftarrow \{0, 1\}^{256}$; $\mathbf{A} \leftarrow R_q^{k \times \ell} := \text{Sam}(\rho)$ 02 $\mathbf{t} \leftarrow R_q^k$ 03 $\mathbf{t}_1 := \text{Power2Round}_q(\mathbf{t}, d)$ 04 $\mathbf{t}_0 := \mathbf{t} - \mathbf{t}_1 \cdot 2^d$ 05 return $pk = (\rho, \mathbf{t}_1, \mathbf{t}_0)$ </pre>
--

Figure 15: The lossy instance generator `LossyGen`.

By observing that the output of `LossyGen` is uniformly random over $R_q^{k \times \ell} \times R_q^k$ and the output of `IGen` in Figure 13 is $(\mathbf{A}, \mathbf{A}\mathbf{s}_1 + \mathbf{s}_2)$ where $\mathbf{A} \leftarrow R_q^{k \times \ell}$ and $(\mathbf{s}_1, \mathbf{s}_2) \leftarrow S_\eta^\ell \times S_\eta^k$, we have that

$$\text{Adv}_{\text{ID}}^{\text{LOSS}}(\mathbf{A}) = \text{Adv}_{k, \ell, D}^{\text{MLWE}}(\mathbf{A}),$$

where D is the uniform distribution over S_η .

Lemma 4.5 *If $4\gamma + 2, 2\gamma' < \sqrt{q/2}$ and $\gamma' < \gamma\beta$, and $\ell \leq k$, then `ID` has ε_{ls} -lossy soundness for*

$$\varepsilon_{\text{ls}} \leq \frac{1}{|\text{ChSet}|} + 2 \cdot |\text{ChSet}|^2 \cdot \left(\frac{32\gamma\gamma'}{q} \right)^{nk}.$$

Our proof follows the framework from [KW03, AFLT12]— first, in Lemma 4.6 we show that if \mathbf{A}, \mathbf{t} are chosen at random, then a particular linear equation is unlikely to have any solutions. Then to prove Lemma 4.5, we show that if \mathbf{C} , who outputs the first message (\mathbf{w}_1, St) in the `LOSSY-IMP` game (see Figure 16) is able to correctly respond to more than one random challenge c , then the previously mentioned linear equation will have a solution, which with high probability is not possible. Therefore we conclude that for virtually all \mathbf{A}, \mathbf{t} output by `LossyGen`, there exists (at most) only one challenge for which the prover can respond to, and therefore his success probability is at most $1/|\text{ChSet}|$.

Lemma 4.6 *Let α_1, α_2 be positive integers less than $\sqrt{q/2}$ and \mathcal{D} be a set of elements in $R \setminus \{0\}$ with coefficients less than $\sqrt{q/2}$. Also, let d be such that $2^d < 2\alpha_1$. Then*

$$\begin{aligned} & \Pr_{\mathbf{A} \leftarrow R_q^{k \times \ell}, \mathbf{t} \leftarrow R_q^k} [\exists (\mathbf{z}_1, \mathbf{z}_2, c) \in S_{\alpha_1}^\ell \times S_{\alpha_2}^k \times \mathcal{D} \text{ s.t. } \mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = \mathbf{c}\mathbf{t}_1 \cdot 2^d] \\ & \leq 2 \cdot |\mathcal{D}| \cdot \left(\frac{(2\alpha_1 + 1)^\ell \cdot (2\alpha_2 + 1)^k}{q^k} \right)^n. \end{aligned} \quad (15)$$

where $\mathbf{t}_1 := \text{Power2Round}_q(\mathbf{t}, d)$.

Proof. Case 1: We will first handle the case where $\mathbf{z}_1 = \mathbf{0}$. In this case Equation (15) becomes

$$\Pr_{\mathbf{t} \leftarrow R_q^k} [\exists (\mathbf{z}_2, c) \in S_{\alpha_2}^k \times \mathcal{D} \text{ such that } \mathbf{z}_2 = \mathbf{c}\mathbf{t}_1 \cdot 2^d].$$

Because $0 < \|c\|_\infty < \sqrt{q/2}$ and $q = 5 \pmod{8}$, we know that c is invertible in R_q ([LN17, Lemma 2.2]). The above probability therefore becomes

$$\begin{aligned} & \Pr_{\mathbf{t} \leftarrow R_q^k} [\exists (\mathbf{z}_2, c) \in S_{\alpha_2}^k \times \mathcal{D} \text{ such that } \mathbf{z}_2 \cdot (2^d c)^{-1} = \mathbf{t}_1] \\ & \leq \sum_{\mathbf{z}_2 \in S_{\alpha_2}^k, c \in \mathcal{D}} \Pr_{\mathbf{t} \leftarrow R_q^k} [\mathbf{t}_1 = \mathbf{z}_2 \cdot (2^d c)^{-1}]. \end{aligned}$$

For $\mathbf{t} \in R_q^k$, the most frequent value of each coefficient of \mathbf{t}_1 occurs at most 2^d times. Thus

$$\begin{aligned} \sum_{\mathbf{z}_2 \in S_{\alpha_2}^k, c \in \mathcal{D}} \Pr_{\mathbf{t} \leftarrow R_q^k} [\mathbf{t}_1 = \mathbf{z}_2 \cdot (2^d c)^{-1}] &\leq \sum_{\mathbf{z}_2 \in S_{\alpha_2}^k, c \in \mathcal{D}} \left(\frac{2^d}{q}\right)^{nk} \\ &= \left(\frac{(2\alpha_2 + 1) \cdot 2^d}{q}\right)^{nk} \cdot |\mathcal{D}|. \end{aligned} \quad (16)$$

Case 2: We now move to the case where $\mathbf{z}_1 \neq \mathbf{0}$. Let $(\mathbf{z}_1 \neq \mathbf{0}, \mathbf{z}_2, c)$ be any triple and assume without loss of generality that the first polynomial in \mathbf{z}_1 is non-zero. We can then write

$$\begin{aligned} &\Pr_{\mathbf{A} \leftarrow R_q^{k \times \ell}, \mathbf{t} \leftarrow R_q^k} [\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}_1 \cdot 2^d] \\ &= \Pr_{\mathbf{a} \leftarrow R_q^k, \mathbf{A}' \leftarrow R_q^{k \times (k-1)}, \mathbf{t} \leftarrow R_q^k} [\mathbf{a}z + \mathbf{A}'\mathbf{z}'_1 + \mathbf{z}_2 = c\mathbf{t}_1 \cdot 2^d] \\ &= \Pr_{\mathbf{a} \leftarrow R_q^k} [\mathbf{a}z = -\mathbf{A}'\mathbf{z}'_1 - \mathbf{z}_2 + c\mathbf{t}_1 \cdot 2^d], \end{aligned}$$

where $\mathbf{z}_1 := \begin{bmatrix} z \\ \mathbf{z}'_1 \end{bmatrix}$. Since $\|z\|_\infty < \sqrt{q/2}$ and $q = 5 \pmod{8}$, we again know that z is invertible in R_q . The above probability therefore becomes

$$\Pr_{\mathbf{a} \leftarrow R_q^k} [\mathbf{a} = z^{-1} \cdot (-\mathbf{A}'\mathbf{z}'_1 - \mathbf{z}_2 + c\mathbf{t}_1 \cdot 2^d)] = q^{-nk}.$$

Thus by the union bound, we can upper-bound Equation (15) when $\mathbf{z}_1 \neq \mathbf{0}$ by

$$\sum_{\mathbf{z}_1 \in S_{\alpha_1}^\ell \setminus \{\mathbf{0}\}, \mathbf{z}_2 \in S_{\alpha_2}^k, c \in \mathcal{D}} q^{-nk} < \left(\frac{(2\alpha_1 + 1)^\ell \cdot (2\alpha_2 + 1)^k}{q^k}\right)^n \cdot |\mathcal{D}|.$$

Combining the above with Equation (16) and using the assumption that $2^d < 2\alpha$, we get the statement in the claim of the Lemma. \square

Proof of Lemma 4.5. Consider an unbounded adversary C that is executed in game LOSSY-IMP of Figure 16.

GAME LOSSY-IMP:
01 $pk_{\text{is}} := (\rho, \mathbf{t}_1, \mathbf{t}_0) \leftarrow \text{LossyGen}(\text{par})$
02 $(\mathbf{w}_1, St) \leftarrow C(pk_{\text{is}})$
03 $c \leftarrow \text{ChSet}$
04 $(\mathbf{z}, \mathbf{h}) \leftarrow C(St, c)$
05 **return** $\llbracket \mathbf{w}_1 = \text{UseHint}_q(\mathbf{h}, \mathbf{A}\mathbf{z} - c\mathbf{t}_1 \cdot 2^d, 2\gamma) \rrbracket$ **and** $\llbracket \|\mathbf{z}\|_\infty < \gamma' - \beta \rrbracket$

Figure 16: The lossy impersonation game LOSSY-IMP.

Suppose that for some \mathbf{w}_1 , there exist two $c \neq c' \in \text{ChSet}$ and two $(\mathbf{z}, \mathbf{h}), (\mathbf{z}', \mathbf{h}')$ that lead to C winning. In other words, $\|\mathbf{z}\|_\infty, \|\mathbf{z}'\|_\infty < \gamma' - \beta$ and

$$\begin{aligned} \mathbf{w}_1 &= \text{UseHint}_q(\mathbf{h}, \mathbf{A}\mathbf{z} - \mathbf{t}_1 c \cdot 2^d, 2\gamma), \\ \mathbf{w}_1 &= \text{UseHint}_q(\mathbf{h}', \mathbf{A}\mathbf{z}' - \mathbf{t}_1 c' \cdot 2^d, 2\gamma). \end{aligned}$$

By Lemma 4.1, we know that the above implies

$$\begin{aligned} \|\mathbf{A}\mathbf{z} - \mathbf{t}_1 c \cdot 2^d - \mathbf{w}_1 \cdot 2\gamma\|_\infty &\leq 2\gamma + 1, \\ \|\mathbf{A}\mathbf{z}' - \mathbf{t}_1 c' \cdot 2^d - \mathbf{w}_1 \cdot 2\gamma\|_\infty &\leq 2\gamma + 1. \end{aligned}$$

By the triangular inequality, this implies that

$$\|\mathbf{A}(\mathbf{z} - \mathbf{z}') - \mathbf{t}_1 \cdot 2^d \cdot (c - c')\|_\infty \leq 4\gamma + 2,$$

which can be rewritten as

$$\mathbf{A}(\mathbf{z} - \mathbf{z}') + \mathbf{u} = \mathbf{t}_1 \cdot 2^d \cdot (c - c') \quad (17)$$

for some \mathbf{u} such that $\|\mathbf{u}\|_\infty \leq 4\gamma + 2$ (and $\|\mathbf{z} - \mathbf{z}'\|_\infty \leq 2(\gamma' - \beta - 1)$).

If $\mathbf{A} \leftarrow R_q^{k \times \ell}$ and $\mathbf{t} \leftarrow R_q^k$, then Lemma 4.6 tells us that Equation (17) is satisfied with probability less than

$$2 \cdot |\text{ChSet}|^2 \cdot \frac{(4(\gamma' - \beta))^{n\ell} \cdot (8\gamma + 5)^{nk}}{q^{nk}} < 2 \cdot |\text{ChSet}|^2 \cdot \left(\frac{32\gamma\gamma'}{q}\right)^{nk}.$$

Thus, except with the above probability, for every \mathbf{w}_1 , there is at most one possible c that allows \mathbf{C} to win. Therefore, except with the above probability, \mathbf{C} has at most a $1/|\text{ChSet}|$ chance of winning. \square

4.3.4 Min Entropy

In Lemma 4.7 we will prove (using exactly the same technique as in Lemma 4.6) that the \mathbf{w}_1 sent by the honest prover in the first step is extremely likely to be distinct for every run of the protocol.

Lemma 4.7 *If $2\gamma, 2\gamma' < \sqrt{q/2}$ and $\ell \leq k$, then the identification scheme ID in Figure 13 has*

$$\alpha > n\ell \cdot \log \left(\min \left\{ \frac{q}{(4\gamma + 1)(4\gamma' + 1)}, 2\gamma' - 1 \right\} \right)$$

bits of min-entropy (as in Definition 2.6).

Proof. We first claim that

$$\begin{aligned} & \Pr_{\mathbf{A} \leftarrow R_q^{k \times \ell}} [\exists \mathbf{y} \neq \mathbf{y}' \in S_{\gamma'-1}^\ell \text{ s.t. } \text{HighBits}_q(\mathbf{A}\mathbf{y}, 2\gamma) = \text{HighBits}_q(\mathbf{A}\mathbf{y}', 2\gamma)] \\ & < \left(\frac{(4\gamma + 1)(4\gamma' + 1)}{q} \right)^{nk}. \end{aligned} \quad (18)$$

From Equation (18), we know that with probability at least $1 - \left(\frac{(4\gamma+1)(4\gamma'+1)}{q}\right)^{nk}$ over the choice of $\mathbf{A} \leftarrow R_q^{k \times \ell}$, each $W = \text{HighBits}_q(\mathbf{A}\mathbf{y}, 2\gamma)$ has exactly a $\frac{1}{|S_{\gamma'-1}^\ell|} = (2\gamma' - 1)^{-n\ell}$ probability of being output.

The claim in the lemma follows directly from Definition 2.6 and the assumption that $k \geq \ell$

It remains to prove Equation (18). If we define

$$\text{Decompose}_q(\mathbf{A}\mathbf{y}, 2\gamma) = (\mathbf{w}_1, \mathbf{w}_0)$$

and

$$\text{Decompose}_q(\mathbf{A}\mathbf{y}', 2\gamma) = (\mathbf{w}'_1, \mathbf{w}'_0),$$

then $\text{HighBits}_q(\mathbf{A}\mathbf{y}, 2\gamma) = \text{HighBits}_q(\mathbf{A}\mathbf{y}', 2\gamma)$ implies that $\mathbf{A}\mathbf{y} = \mathbf{w}_1 \cdot 2\gamma + \mathbf{w}_0$ and $\mathbf{A}\mathbf{y}' = \mathbf{w}'_1 \cdot 2\gamma + \mathbf{w}'_0$ with $\mathbf{w}_1 = \mathbf{w}'_1$ and $\|\mathbf{w}_0\|_\infty, \|\mathbf{w}'_0\|_\infty \leq \gamma$. Therefore we have

$$\mathbf{A}(\mathbf{y} - \mathbf{y}') - (\mathbf{w}_0 - \mathbf{w}'_0) = \mathbf{0} \quad (19)$$

with

$$\|\mathbf{y} - \mathbf{y}'\|_\infty < 2\gamma', \|\mathbf{w}_0 - \mathbf{w}'_0\|_\infty \leq 2\gamma.$$

Since $2\gamma, 2\gamma' < \sqrt{q/2}$, the same argument as in Case 2 of the proof of Lemma 4.6 shows that the probability over the choice of $\mathbf{A} \leftarrow R_q^{k \times \ell}$, that there exist two non-zero elements of norm less than 2γ and $2\gamma'$, respectively, that satisfy Equation (19) is at most

$$\left(\frac{(4\gamma + 1)^\ell (4\gamma' + 1)^k}{q^k} \right)^n \leq \left(\frac{(4\gamma + 1)(4\gamma' + 1)}{q} \right)^{nk}.$$

This proves Equation (18). \square

4.3.5 Computational Unique Response

In this section we prove that our scheme satisfies the Computational Unique Response (CUR) property required for strong-unforgeability of the signature scheme.

Lemma 4.8 *If $4\gamma + 2, 2\gamma' < \sqrt{q/2}$ and $\gamma' < \gamma\beta$, and $\ell \leq k$ (i.e. the same conditions as in Lemma 4.5), then $\text{Adv}_{\text{ID}}^{\text{CUR}}(\mathbf{A}) < \left(\frac{32\gamma\gamma'}{q}\right)^{nk}$ for every (even unbounded) adversary \mathbf{A} .*

Proof. Let $(W, c, Z) = (\mathbf{w}_1, c, (\mathbf{z}, \mathbf{h}))$ be any valid transcript. If \mathbf{A} is able to generate a valid $Z' = (\mathbf{h}', \mathbf{z}') \neq Z$ such that $\mathbf{V}(pk = (\mathbf{A}, \mathbf{t}_1), \mathbf{w}_1, c, (\mathbf{z}', \mathbf{h}')) = 1$, then Lemma 4.9 implies that there exist \mathbf{v}, \mathbf{u} with $\|\mathbf{v}\|_\infty < 2(\gamma' - \beta)$, $\|\mathbf{u}\|_\infty \leq 4\gamma + 2$ such that $\mathbf{A}\mathbf{v} + \mathbf{u} = \mathbf{0}$. Since $2(\gamma' - \beta), 4\gamma + 2 < \sqrt{q/2}$, the same argument as in Case 2 of the proof of Lemma 4.6 shows that the probability over the choice of $\mathbf{A} \leftarrow R_q^{k \times \ell}$, that there exist such \mathbf{v}, \mathbf{u} is at most

$$\frac{(4(\gamma' - \beta))^{n\ell} \cdot (8\gamma + 5)^{nk}}{q^{nk}} < \left(\frac{32\gamma\gamma'}{q}\right)^{nk}.$$

□

Lemma 4.9 *If $(\mathbf{w}_1, c, (\mathbf{z}, \mathbf{h}))$ and $(\mathbf{w}_1, c, (\mathbf{z}', \mathbf{h}'))$ are such that $\mathbf{V}(pk, \mathbf{w}_1, c, (\mathbf{z}, \mathbf{h})) = \mathbf{V}(pk, \mathbf{w}_1, c, (\mathbf{z}', \mathbf{h}')) = 1$ and $(\mathbf{z}, \mathbf{h}) \neq (\mathbf{z}', \mathbf{h}')$, then there exist \mathbf{v}, \mathbf{u} such that $\|\mathbf{v}\|_\infty < 2(\gamma' - \beta)$, $\|\mathbf{u}\|_\infty \leq 4\gamma + 2$ such that $\mathbf{A}\mathbf{v} + \mathbf{u} = \mathbf{0}$.*

Proof. The two conditions of the Lemma imply that

$$\mathbf{w}_1 = \text{UseHint}_q(\mathbf{h}, \mathbf{A}\mathbf{z} - c\mathbf{t}_1 \cdot 2^d, 2\gamma),$$

$$\mathbf{w}_1 = \text{UseHint}_q(\mathbf{h}', \mathbf{A}\mathbf{z}' - c\mathbf{t}_1 \cdot 2^d, 2\gamma).$$

We first point out that it must be that $\mathbf{z} \neq \mathbf{z}'$. This is because Lemma 4.1 implies that if $\mathbf{z} = \mathbf{z}'$ then necessarily $\mathbf{h} = \mathbf{h}'$ (and then $Z = Z'$). The above two equations imply (again by Lemma 4.1) that

$$\|\mathbf{A}\mathbf{z} - c\mathbf{t}_1 \cdot 2^d - \mathbf{w}_1 \cdot 2\gamma\|_\infty \leq 2\gamma + 1,$$

$$\|\mathbf{A}\mathbf{z}' - c\mathbf{t}_1 \cdot 2^d - \mathbf{w}_1 \cdot 2\gamma\|_\infty \leq 2\gamma + 1.$$

By the triangular inequality, this can be rewritten as

$$\mathbf{A}(\mathbf{z} - \mathbf{z}') + \mathbf{u} = \mathbf{0}$$

for some \mathbf{u} such that $\|\mathbf{u}\| \leq 4\gamma_2 + 2$ and $\|\mathbf{z} - \mathbf{z}'\| < 2(\gamma' - \beta)$. □

4.4 The Dilithium-QROM Signature Scheme and Concrete Parameters

In this section, we describe the signature scheme Dilithium-QROM (Figure 17) which is obtained via the Fiat-Shamir transform from the scheme ID of Figure 13 and using $\text{Sam}(K \parallel \cdot)$ as a pseudorandom function. We then instantiate it with concrete parameters (Table 1) and compare them for the same security level with those in [DLL⁺17].

The parameters for our scheme are dictated by the requirements for the scheme to be strongly-unforgeable in Theorem 3.1 which gives an upper bound on $\text{Adv}_{\text{Dilithium-QROM}}^{\text{SUF-CMA}}(\mathbf{A})$. Following [KMP16], for “ κ bits of quantum security” for Dilithium-QROM we require that for all quantum adversaries \mathbf{A} running in time at most 2^κ ,

$$\text{Adv}_{\text{Dilithium-QROM}}^{\text{SUF-CMA}}(\mathbf{A}) / \text{Time}(\mathbf{A}) \leq 2^{-\kappa}. \quad (20)$$

To this end, we need to put bounds on the parameters $\varepsilon_{\text{ls}}, \varepsilon_{\text{zk}}$, and α . Lemma 4.3 tells us that

$$\varepsilon_{\text{zk}} = 0.$$

To lower-bound α , note that in the parameters, we always have $2\gamma = 2\gamma' < \sqrt{q/2}$, and therefore we can apply Lemma 4.7 and obtain that α is greater than 2900. Thus the $2^{-\alpha}$ term has absolutely no practical effect in Theorem 3.1 for the parameters in Section 4.4.

<pre> Sign((sk, K), M) 01 $\kappa := 0$ 02 $\mathbf{A} \leftarrow R_q^{k \times \ell} := \text{Sam}(\rho)$ 03 while $(\mathbf{z}, \mathbf{h}) = \perp$ and $\kappa \leq 200/(1 - \delta)$ do 04 $\kappa := \kappa + 1$ 05 $\mathbf{y} \leftarrow S_{\gamma'}^\ell := \text{Sam}(K \parallel M \parallel \kappa)$ 06 $\mathbf{w} := \mathbf{A}\mathbf{y}$ 07 $\mathbf{w}_1 := \text{HighBits}_q(\mathbf{w}, 2\gamma)$ 08 $c := \text{H}(\mathbf{w}_1 \parallel M)$ 09 $\mathbf{z} := \mathbf{y} + c\mathbf{s}_1$ 10 if $\ \mathbf{z}\ _\infty \geq \gamma' - \beta$ or $\ \text{LowBits}_q(\mathbf{w} - c\mathbf{s}_2, 2\gamma)\ _\infty \geq \gamma - \beta$ then $(\mathbf{z}, \mathbf{h}) := \perp$ 11 else $\mathbf{h} := \text{MakeHint}_q(-c\mathbf{t}_0, \mathbf{w} - c\mathbf{s}_2 + c\mathbf{t}_0, 2\gamma)$ 12 return $\sigma = (\mathbf{z}, \mathbf{h}, c)$ Ver(pk, M, $\sigma = (\mathbf{z}, \mathbf{h}, c)$) 13 $\mathbf{A} \leftarrow R_q^{k \times \ell} := \text{Sam}(\rho)$ 14 $\mathbf{w}'_1 := \text{UseHint}_q(\mathbf{h}, \mathbf{A}\mathbf{z} - c\mathbf{t}_1 \cdot 2^d, 2\gamma)$ 15 return $[\ \mathbf{z}\ _\infty < \gamma' - \beta]$ and $[c = \text{H}(\mathbf{w}'_1 \parallel M)]$ </pre>

Figure 17: Our signature scheme Dilithium-QROM := DFS[ID]. The key generation algorithm is IGen from Figure 13, where the secret key also contains a random key K for the pseudorandom function $\text{Sam}(K \parallel \cdot)$. The bound $200/(1 - \delta)$ on κ can be ignored as there is only a $\delta^{200/(1-\delta)} < \exp(-200)$ chance that it will be reached in any call to the signing procedure. Its presence is for consistency with the generic signing algorithm in Section 3.1.

Lemma 4.8 states that as long as $4\gamma + 2$ and $2\gamma' < \sqrt{q/2}$, we will have $\text{Adv}_{\text{ID}}^{\text{CUR}}(\text{C}) < \left(\frac{32\gamma\gamma'}{q}\right)^{nk}$. The parameters in Table 1 indeed satisfy the preconditions, and so $\text{Adv}_{\text{ID}}^{\text{CUR}}(\text{C}) < \left(\frac{32\gamma\gamma'}{q}\right)^{nk} < 2^{-865}$.

We finally turn to bounding ε_{ls} . Notice that Lemma 4.5 directly implies that

$$\varepsilon_{\text{ls}} \leq \frac{1}{|\text{ChSet}|} + 2 \cdot |\text{ChSet}|^2 \cdot \left(\frac{32\gamma\gamma'}{q}\right)^{nk}.$$

The size of the challenge set ChSet defined in Equation (8) is larger than 2^{265} , and so the above is at most

$$\varepsilon_{\text{ls}} \leq 2^{-265} + 2^{-334} \leq 2^{-264}.$$

Plugging everything into the equation at the end of Section 3.1, we obtain

$$\begin{aligned} \text{Adv}_{\text{Dilithium-QROM}}^{\text{UF-CMA}}(\text{A}) &\leq \text{Adv}_{\text{ID}}^{\text{LOSS}}(\text{B}) + \text{Adv}_{\text{ID}}^{\text{CUR}}(\text{C}) + 8 \cdot (Q_{\text{H}} + 1)^2 \cdot \varepsilon_{\text{ls}} \\ &\quad + \text{Adv}_{\text{Sam}}^{\text{PR}}(\text{D}) + \frac{200}{(1 - \delta)} \cdot Q_{\text{S}} \cdot \varepsilon_{\text{zk}} + 2^{-\alpha} \\ &< \text{Adv}_{\text{ID}}^{\text{MLWE}}(\text{B}) + Q_{\text{H}}^2 \cdot 2^{-261} + \text{Adv}_{\text{Sam}}^{\text{PR}}(\text{D}). \end{aligned}$$

Table 1 also shows that the parameters of the MLWE problem are chosen such that it provides 128 bits of quantum security (using the same metric as was used in the original Dilithium scheme [DLL+17]). Assuming Sam provides 128 bits security when used as a pseudorandom function, we conclude that for all quantum adversaries running in time at most 2^{128} and making $1 \leq Q_{\text{H}} \leq 2^{128}$ (quantum) queries to H, and we have

$$\frac{\text{Adv}_{\text{Dilithium-QROM}}^{\text{UF-CMA}}(\text{A})}{\text{Time}(\text{A})} \leq \frac{\text{Adv}_{\text{ID}}^{\text{MLWE}}(\text{B})}{\text{Time}(\text{B})} + \frac{\text{Adv}_{\text{Sam}}^{\text{PR}}(\text{D})}{\text{Time}(\text{D})} + Q_{\text{H}} \cdot 2^{-261} \leq 2^{-128}$$

	Dilithium-QROM		Dilithium [DLL ⁺ 17]	
	recomm.	very high	recomm.	very high
q (ring modulus)	$2^{45} - 21283$	$2^{45} - 21283$	$2^{23} - 8191$	$2^{23} - 8191$
n (ring dimension)	512	512	256	256
(k, ℓ) (dimension of matrix \mathbf{A})	(4, 4)	(5, 5)	(5, 4)	(6, 5)
d (dropped bits from \mathbf{t})	15	15	14^4	14
# of ± 1 's in $c \in \text{ChSet}$	46	46	60	60
γ s.t. $2\gamma \mid q - 1$	905679	905679	261888	261888
γ' (\approx max. sig. coefficient)	905679	905679	523776	523776
η (maximum coefficient of $\mathbf{s}_1, \mathbf{s}_2$)	7	3	5	3
β ($= \eta \cdot (\# \text{ of } \pm 1\text{'s in } c)$)	322	138	275^5	175
pk size (bytes)	7712	9632	1472	1760
sig size (bytes)	5690	7098	2701	3366
Exp. Repeats ($1/(1 - \delta)$ from Lemma 4.4)	4.3	2.2	6.6	4.3
BKZ block-size to break LWE	480	600	485	595
Best Known Classical bit-cost	140	175	141	174
Best Known Quantum bit-cost	127	159	128	158
BKZ block-size to break SIS	NA	NA	475	605
Best Known Classical bit-cost	NA	NA	138	176
Best Known Quantum bit-cost	NA	NA	125	160

Table 1: Parameters for Dilithium-QROM and Dilithium. The security analysis for the MLWE and MSIS problems is as described in [DLL⁺17].

The signature size in Dilithium-QROM is $(n \cdot \ell \cdot (\lceil \log(2\gamma) \rceil) + nk + 46 \cdot (\log(n) + 1))/8$ bytes, while the public key is $(n \cdot k \cdot (\lceil \log(q) \rceil - d) + 256)/8$ bytes.

In Table 1, we compare the parameters from the current scheme, which can be proved secure based on the hardness of MLWE in the QROM, to those of the original Dilithium scheme from [DLL⁺17], which only has a classical security reduction from the combination of MLWE and MSIS (we introduce this latter problem in the next section). One can see that the sum of the public key and signature sizes are approximately 3.2 times larger in Dilithium-QROM than in Dilithium.

4.5 Security Assumptions for Non-Lossy Schemes

The reduction from the MLWE problem to the hardness of the Dilithium-QROM scheme was a direct consequence of Theorem 3.1, which is itself a combination of Theorem 3.2 and Theorem 3.4. In this section, we consider the security of schemes for which Theorem 3.4 is inapplicable. In particular, in these schemes it is no longer true that a computationally-unbounded adversary cannot win the LOSSY-IMP game. The reason that one would like to use schemes constructed in such a manner is because they turn out to be more efficient. In particular, the original Dilithium scheme⁶ [DLL⁺17], which is virtually identical to the Dilithium-QROM presented in this paper except for the parameter sizes, has outputs (of the public key plus signature) that are smaller by a factor of a little over 3 (see Table 1).

But while the Dilithium scheme has a security reduction from standard lattice problems in the *classical* random-oracle model, there is no such reduction in the quantum random-oracle model. Nevertheless, it is

⁵For added compactness of the public key, the size of d (i.e. the amount of bits that one can drop from \mathbf{t}) can be such that the necessary condition $\|\mathbf{ct}_0\|_\infty < \gamma$ is not always satisfied. This would invalidate the correctness of the scheme – in particular Equation (14). Nevertheless, if this condition is satisfied most of the time and the signer simply checks whether $\|\mathbf{ct}_0\|_\infty < \gamma$ before sending the signature (and aborts the signing attempt otherwise), then the scheme retains its correctness property. Since for security, we assumed that \mathbf{t}_0 is known to the adversary, this check does not affect security. In the Dilithium scheme, this check is performed at the end of the while loop of the signing algorithm.

⁵The β values for Dilithium were chosen such that $\Pr_{\Gamma_{s \leftarrow S_\gamma, c \leftarrow \text{ChSet}}}[\|sc\|_\infty > \beta]$ is very close to 0. Increasing / decreasing the value of β changes the value δ , which has an effect on the run-time of the scheme.

⁶We refer to the deterministic version of the scheme.

unclear whether this lack of reduction implies any weakness against quantum attacks. It would therefore be useful to understand exactly what assumptions the more efficient scheme is relying on in the quantum random-oracle model.

Let us suppose that the parameters for the Dilithium scheme are set such that Theorem 3.2 is still applicable. That is, suppose that $\varepsilon_{\text{zk}} = 0$, α is very large, and the scheme is commitment-recoverable. In this case, ignoring the $2^{-\alpha+1}$ term, Theorem 3.2 states that the security of the full signature scheme is exactly the security of the UF-NMA signature scheme in the quantum random-oracle model. Since the adversary does not obtain any valid signatures in the UF-NMA security game, the security assumption of such signatures is non-interactive.

Below, we recall the standard MSIS assumption and then define a new assumption, SelfTargetMSIS, upon which the security of Dilithium is based. We also point out that in the *classical* random-oracle model, there is a (non-tight) reduction from the MSIS to the SelfTargetMSIS problem. Then we show that the Dilithium scheme for which Theorem 3.4 is not necessarily applicable, still has a security reduction from the combination of MLWE and SelfTargetMSIS problems.

4.5.1 The MSIS and SelfTargetMSIS Problems

The MSIS problem [LS15] is a generalization of the SIS [Ajt96] and Ring-SIS [PR06, LM06] problems in the same way that MLWE is a generalization of LWE and Ring-LWE. To an algorithm A we associate the advantage function $\text{Adv}_{m,k,\gamma}^{\text{MSIS}}(A)$ to solve the (Hermite Normal Form) MSIS $_{m,k,\gamma}$ problem over the ring R_q as

$$\text{Adv}_{m,k,\gamma}^{\text{MSIS}}(A) := \Pr [0 < \|\mathbf{y}\|_\infty \leq \gamma \wedge [\mathbf{I} \mid \mathbf{A}] \cdot \mathbf{y} = \mathbf{0} \mid \mathbf{A} \leftarrow R_q^{m \times k}; \mathbf{y} \leftarrow \mathbf{A}(\mathbf{A})].$$

As for SIS and Ring-SIS, it was shown that solving MSIS for certain parameters is as hard as worst-case instances of lattice problems over algebraic lattices of a certain form [LS15].

Suppose that $H : \{0, 1\}^* \rightarrow \text{ChSet}$ is a cryptographic hash function. To an algorithm A we associate the advantage function $\text{Adv}_{H,m,k,\gamma}^{\text{SelfTargetMSIS}}(A)$ to solve the SelfTargetMSIS $_{H,m,k,\gamma}$ problem over the ring R_q as

$$\text{Adv}_{H,m,k,\gamma}^{\text{SelfTargetMSIS}}(A) :=$$

$$\Pr \left[\begin{array}{l} \|\mathbf{y}\|_\infty \leq \gamma \\ \wedge H([\mathbf{I} \mid \mathbf{A}] \cdot \mathbf{y} \parallel M) = c \end{array} \mid \mathbf{A} \leftarrow R_q^{m \times k}; \left(\mathbf{y} := \begin{bmatrix} \mathbf{r} \\ c \end{bmatrix}, M \right) \leftarrow \mathbf{A}^{|\mathbf{H}}(\mathbf{A}) \right].$$

If A only has classical access to H , then there is a reduction, using the forking lemma [PS00, BN06], to prove that $\text{Adv}_{H,m,k,\gamma}^{\text{SelfTargetMSIS}}(B) \approx \sqrt{\text{Adv}_{m,k,2\gamma}^{\text{MSIS}}(A)/Q_H}$, where Q_H is the number of classical queries to H .⁷ This reduction is standard and is implicit in the (classical) security proofs of digital signatures based on the hardness of the SIS problem (cf. [Lyu12, DLL⁺17]). For completeness, we give its sketch below.

If A is a solver for the SelfTargetMSIS $_{H,m,k,\gamma}$ problem, then B passes the \mathbf{A} from his MSIS $_{m,k,2\gamma}$ instance to A and replies to A 's queries $H(W \parallel M)$ with uniformly random $c \in \text{ChSet}$. If A returns a solution $\left(\mathbf{y} = \begin{bmatrix} \mathbf{r} \\ c \end{bmatrix}, M \right)$ to SelfTargetMSIS $_{H,m,k,\gamma}$, then B reruns A with the same randomness, but reprograms the “winning” query $H([\mathbf{I} \mid \mathbf{A}] \cdot \mathbf{y} \parallel M)$ to a different random element c' . (We’re assuming that in order to “win” with (\mathbf{y}, M) , A must have queried $H([\mathbf{I} \mid \mathbf{A}] \cdot \mathbf{y} \parallel M)$ – if not then he is able to break the second-preimage resistance of H .) The forking lemma states that then A has a $1/Q_H$ probability of succeeding using the same query (see, e.g., [BN06]). If A then outputs a solution $\left(\mathbf{y}' = \begin{bmatrix} \mathbf{r}' \\ c' \end{bmatrix}, M \right)$ such that $H([\mathbf{I} \mid \mathbf{A}] \cdot \mathbf{y}' \parallel M) = c'$ with $[\mathbf{I} \mid \mathbf{A}] \cdot \mathbf{y} = [\mathbf{I} \mid \mathbf{A}] \cdot \mathbf{y}'$, then $[\mathbf{I} \mid \mathbf{A}] \cdot (\mathbf{y} - \mathbf{y}') = \mathbf{0}$ with $\|\mathbf{y}\|_\infty, \|\mathbf{y}'\|_\infty \leq \gamma$. Since $c \neq c'$, we know that $0 < \|\mathbf{y} - \mathbf{y}'\|_\infty \leq 2\gamma$, and B has a solution to MSIS $_{m,k,2\gamma}$.

Note that the reduction loses a factor of Q_H in the success probability and also the size of the coefficients in the solution to MSIS is larger than that in SelfTargetMSIS, which could lead one to conclude that SelfTargetMSIS is an easier problem. But this is the standard “tightness loss” in the Fiat-Shamir transform protocol that is not constructed from lossy identification schemes. Since its first use over three decades ago, however, there have been no attacks that break the Fiat-Shamir signature with more success than the underlying problem upon which the scheme is based.

⁷This can be improved to $Q_H \text{Adv}_{H,m,k,\gamma}^{\text{SelfTargetMSIS}}(B)/\text{Time}(B) \approx \text{Adv}_{m,k,2\gamma}^{\text{MSIS}}(A)/\text{Time}(A)$.

It is therefore very reasonable to assume that unless there is some algebraic relationship between \mathbf{H} and \mathbf{A} , then the $\text{SelfTargetMSIS}_{\mathbf{H},m,k,\gamma}$ problem is as hard as $\text{MSIS}_{m,k,2\gamma}$ problem.⁸ If \mathbf{A} has quantum access to \mathbf{H} , then one can no longer use the forking lemma (which is essentially rewinding) in order to get a non-tight equivalence between the MSIS and SelfTargetMSIS problems. Nevertheless, as long as \mathbf{H} is second-preimage resistant against quantum attackers, it would again appear that some relationship between \mathbf{A} and \mathbf{H} needs to be exploited by the quantum attacker in order to solve the SelfTargetMSIS problem more efficiently than the MSIS one.

4.5.2 Security based on MLWE, MSIS, and SelfTargetMSIS in the QROM

The QROM security of (deterministic) Dilithium can be expressed as

$$\text{Adv}_{\text{Dilithium}}^{\text{sUF-CMA}}(\mathbf{A}) \leq \text{Adv}_{k,\ell,D}^{\text{MLWE}}(\mathbf{B}) + \text{Adv}_{\mathbf{H},k,\ell+1,\zeta}^{\text{SelfTargetMSIS}}(\mathbf{C}) + \text{Adv}_{\text{Sam}}^{\text{PR}}(\mathbf{D}) + \text{Adv}_{k,\ell,\zeta'}^{\text{MSIS}}(\mathbf{E}) + 2^{-\alpha+1}, \quad (21)$$

for D a uniform distribution over S_η ,

$$\zeta = \max\{\gamma' - \beta, 2\gamma + 1 + 2^{d-1} \cdot \rho\}, \quad (22)$$

where ρ is the number of ± 1 's in the challenge set ChSet , and

$$\zeta' = \max\{2(\gamma' - \beta), 4\gamma + 2\}. \quad (23)$$

For a proof that the min-entropy α is greater than 255 for both parameter sets, see Appendix C. For strong unforgeability, Lemma 4.9 directly implies that $\text{Adv}_{\text{ID}}^{\text{CUR}}(\mathbf{A}) \leq \text{Adv}_{k,\ell,\zeta'}^{\text{MSIS}}(\mathbf{E})$. The bound in Equation (21) is then obtained by combining Theorem 3.2 with the computations from Section 4.3 and Lemma 4.10.

Lemma 4.10 *For any quantum adversary \mathbf{A} against UF-NMA security that issues at most $Q_{\mathbf{H}}$ queries to the quantum random oracle $|\mathbf{H}\rangle$, there exist quantum adversaries \mathbf{B} and \mathbf{C} such that*

$$\text{Adv}_{\text{SIG}}^{\text{UF-NMA}}(\mathbf{A}) \leq \text{Adv}_{k,\ell,D}^{\text{MLWE}}(\mathbf{B}) + \text{Adv}_{\mathbf{H},k,\ell+1,\zeta}^{\text{SelfTargetMSIS}}(\mathbf{C}), \quad (24)$$

and $\text{Time}(\mathbf{B}) = \text{Time}(\mathbf{C}) = \text{Time}(\mathbf{A}) + Q_{\mathbf{H}}$. D is the uniform distribution over S_η and ζ, ρ are as in Equation (22).

Proof. Given an $\mathbf{A}' = [\mathbf{I} \mid \mathbf{A}']$ for $\mathbf{A}'' \in R_q^{k \times (\ell+1)}$, \mathbf{C} writes $\mathbf{A}'' = [\mathbf{A} \mid \mathbf{t}]$ and sets (\mathbf{A}, \mathbf{t}) as the public key of the signature scheme and sends it to \mathbf{A} . If the pk generated by IGen is indistinguishable from uniform over $R_q^{k \times \ell} \times R_q^k$ (i.e. if the $\text{MLWE}_{k,\ell,D}$ problem is hard), then with probability $\text{Adv}_{\mathbf{H},k,\ell+1,\gamma}^{\text{SelfTargetMSIS}}(\mathbf{A})$, \mathbf{A} will return a signature $(c, (\mathbf{z}, \mathbf{h}))$ of some message M such that $\|\mathbf{z}\|_\infty < \gamma' - \beta$ satisfying the verification equation

$$c = \text{H}(\text{UseHint}_q(\mathbf{h}, \mathbf{Az} - \mathbf{ct}_1 \cdot 2^d, 2\gamma) \parallel M).$$

From Lemma 4.1, we know that the above equality can be rewritten as

$$c = \text{H}(\mathbf{Az} - \mathbf{ct}_1 \cdot 2^d + \mathbf{u} \parallel M),$$

where $\|\mathbf{u}\|_\infty \leq 2\gamma + 1$. Using the fact that $\mathbf{t} = \mathbf{t}_1 \cdot 2^d + \mathbf{t}_0$ where $\|\mathbf{t}_0\|_\infty \leq 2^{d-1}$, the above can be again re-written as

$$c = \text{H}(\mathbf{Az} - \mathbf{ct} + (\mathbf{ct}_0 + \mathbf{u}) \parallel M) = \text{H}(\mathbf{Az} - \mathbf{ct} + \mathbf{u}' \parallel M), \quad (25)$$

where $\|\mathbf{u}'\|_\infty \leq \|\mathbf{u}\|_\infty + \|\mathbf{ct}_0\|_\infty \leq 2\gamma + 1 + 2^{d-1} \cdot \rho$. Equation (25), gives us a $\mathbf{y} \in R_q^{k \times (k+\ell+1)}$ such that $\text{H}(\mathbf{A}'\mathbf{y} \parallel M) = c$ where $\mathbf{y} = \begin{bmatrix} \mathbf{r} \\ c \end{bmatrix}$ and $\|\mathbf{y}\|_\infty = \|\mathbf{r}\|_\infty = \max\{\gamma' - \beta, 2\gamma + 1 + 2^{d-1} \cdot \rho\}$. \square

⁸It's also not clear whether the γ in the SelfTargetMSIS problem actually becomes 2γ in the MSIS problem, or if that's also an artifact of the reduction.

Acknowledgments

Eike Kiltz was supported in part by ERC Project ERCC (FP7/615074) and by DFG SPP 1736 Big Data. Vadim Lyubashevsky was supported by the SNSF ERC Transfer Starting Grant CRETP2-166734-FELICITY and the H2020 Project SAFEcrypto. Christian Schaffner was supported by a NWO VIDI grant (639.022.519). The authors are grateful to Dominique Unruh and the anonymous reviewers for comments and discussions.

References

- [AABN02] Michel Abdalla, Jee Hea An, Mihir Bellare, and Chathuranga Namprempe. From identification to signatures via the Fiat-Shamir transform: Minimizing assumptions for security and forward-security. In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 418–433. Springer, Heidelberg, April / May 2002. (Cited on page 1.)
- [ABB⁺17] Erdem Alkim, Nina Bindel, Johannes A. Buchmann, Özgür Dagdelen, Edward Eaton, Gus Gutoski, Juliane Krämer, and Filip Pawlega. Revisiting TESLA in the quantum random oracle model. In *Post-Quantum Cryptography - 8th International Workshop, PQCrypto 2017, Utrecht, The Netherlands, June 26-28, 2017, Proceedings*, pages 143–162, 2017. (Cited on page 2, 3, 4, 5.)
- [AFLT12] Michel Abdalla, Pierre-Alain Fouque, Vadim Lyubashevsky, and Mehdi Tibouchi. Tightly-secure signatures from lossy identification schemes. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 572–590. Springer, Heidelberg, April 2012. (Cited on page 2, 3, 4, 5, 9, 21, 36.)
- [Ajt96] Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *28th ACM STOC*, pages 99–108. ACM Press, May 1996. (Cited on page 27.)
- [ARU14] Andris Ambainis, Ansis Rosmanis, and Dominique Unruh. Quantum attacks on classical proof systems: The hardness of quantum rewinding. In *55th FOCS*, pages 474–483. IEEE Computer Society Press, October 2014. (Cited on page 3.)
- [BBBV97] Charles H. Bennett, Ethan Bernstein, Gilles Brassard, and Umesh V. Vazirani. Strengths and weaknesses of quantum computing. *SIAM J. Comput.*, 26(5):1510–1523, 1997. (Cited on page 2, 3.)
- [BBC⁺98] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quantum lower bounds by polynomials. In *39th FOCS*, pages 352–361. IEEE Computer Society Press, November 1998. (Cited on page 6.)
- [BDF⁺11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 41–69. Springer, Heidelberg, December 2011. (Cited on page 2, 3, 6, 7.)
- [BG14] Shi Bai and Steven D. Galbraith. An improved compression technique for signatures based on learning with errors. In Josh Benaloh, editor, *CT-RSA 2014*, volume 8366 of *LNCS*, pages 28–47. Springer, Heidelberg, February 2014. (Cited on page 2, 5.)
- [BHLY16] Leon Groot Bruinderink, Andreas Hülsing, Tanja Lange, and Yuval Yarom. Flush, gauss, and reload - A cache attack on the BLISS lattice-based signature scheme. In Benedikt Gierlich and Axel Y. Poschmann, editors, *CHES 2016*, volume 9813 of *LNCS*, pages 323–345. Springer, Heidelberg, August 2016. (Cited on page 3.)
- [BN06] Mihir Bellare and Gregory Neven. Multi-signatures in the plain public-key model and a general forking lemma. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 06*, pages 390–399. ACM Press, October / November 2006. (Cited on page 2, 27.)

- [BPS16] Mihir Bellare, Bertram Poettering, and Douglas Stebila. From identification to signatures, tightly: A framework and generic transforms. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 435–464. Springer, Heidelberg, December 2016. (Cited on page 4, 10.)
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In V. Ashby, editor, *ACM CCS 93*, pages 62–73. ACM Press, November 1993. (Cited on page 2, 7.)
- [BR06] Mihir Bellare and Phillip Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 409–426. Springer, Heidelberg, May / June 2006. (Cited on page 6.)
- [DDLL13] Léo Ducas, Alain Durmus, Tancrede Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal Gaussians. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 40–56. Springer, Heidelberg, August 2013. (Cited on page 2, 3, 5.)
- [DLL⁺17] Léo Ducas, Tancrede Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS - Dilithium: Digital signatures from module lattices. *IACR Cryptology ePrint Archive*, 2017:633, 2017. To appear in TCHES 2018. (Cited on page 2, 3, 4, 16, 17, 24, 25, 26, 27.)
- [DLP14] Léo Ducas, Vadim Lyubashevsky, and Thomas Prest. Efficient identity-based encryption over NTRU lattices. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 22–41. Springer, Heidelberg, December 2014. (Cited on page 2, 3.)
- [EFGT17] Thomas Espitau, Pierre-Alain Fouque, Benoît Gérard, and Mehdi Tibouchi. Side-channel attacks on BLISS lattice-based signatures - exploiting branch tracing against strongswan and electromagnetic emanations in microcontrollers. *IACR Cryptology ePrint Archive*, 2017:505, 2017. (Cited on page 3.)
- [ES15] Edward Eaton and Fang Song. Making existential-unforgeable signatures strongly unforgeable in the quantum random-oracle model. In *10th Conference on the Theory of Quantum Computation, Communication and Cryptography, TQC 2015, May 20-22, 2015, Brussels, Belgium*, pages 147–162, 2015. (Cited on page 2.)
- [FS87] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *CRYPTO'86*, volume 263 of *LNCS*, pages 186–194. Springer, Heidelberg, August 1987. (Cited on page 1.)
- [GLP12] Tim Güneysu, Vadim Lyubashevsky, and Thomas Pöppelmann. Practical lattice-based cryptography: A signature scheme for embedded systems. In Emmanuel Prouff and Patrick Schaumont, editors, *CHES 2012*, volume 7428 of *LNCS*, pages 530–547. Springer, Heidelberg, September 2012. (Cited on page 2, 5.)
- [HRS16] Andreas Hülsing, Joost Rijneveld, and Fang Song. Mitigating multi-target attacks in hash-based signatures. In Chen-Mou Cheng, Kai-Min Chung, Giuseppe Persiano, and Bo-Yin Yang, editors, *PKC 2016, Part I*, volume 9614 of *LNCS*, pages 387–416. Springer, Heidelberg, March 2016. (Cited on page 4, 7, 32.)
- [KMP16] Eike Kiltz, Daniel Masny, and Jiaxin Pan. Optimal security proofs for signatures from identification schemes. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 33–61. Springer, Heidelberg, August 2016. (Cited on page 2, 24.)
- [KTX08] Akinori Kawachi, Keisuke Tanaka, and Keita Xagawa. Concurrently secure identification schemes based on the worst-case hardness of lattice problems. In Josef Pieprzyk, editor, *ASIACRYPT 2008*, volume 5350 of *LNCS*, pages 372–389. Springer, Heidelberg, December 2008. (Cited on page 2, 5.)

- [KW03] Jonathan Katz and Nan Wang. Efficiency improvements for signature schemes with tight security reductions. In Sushil Jajodia, Vijayalakshmi Atluri, and Trent Jaeger, editors, *ACM CCS 03*, pages 155–164. ACM Press, October 2003. (Cited on page 2, 21, 36.)
- [LM06] Vadim Lyubashevsky and Daniele Micciancio. Generalized compact Knapsacks are collision resistant. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *ICALP 2006, Part II*, volume 4052 of *LNCS*, pages 144–155. Springer, Heidelberg, July 2006. (Cited on page 27.)
- [LN17] Vadim Lyubashevsky and Gregory Neven. One-shot verifiable encryption from lattices. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part I*, volume 10210 of *LNCS*, pages 293–323. Springer, Heidelberg, May 2017. (Cited on page 5, 16, 21.)
- [LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 1–23. Springer, Heidelberg, May 2010. (Cited on page 17.)
- [LS15] Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Des. Codes Cryptography*, 75(3):565–599, 2015. (Cited on page 17, 27.)
- [Lyu08] Vadim Lyubashevsky. Lattice-based identification schemes secure under active attacks. In Ronald Cramer, editor, *PKC 2008*, volume 4939 of *LNCS*, pages 162–179. Springer, Heidelberg, March 2008. (Cited on page 2.)
- [Lyu09] Vadim Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 598–616. Springer, Heidelberg, December 2009. (Cited on page 2, 3, 4, 5, 10, 16.)
- [Lyu12] Vadim Lyubashevsky. Lattice signatures without trapdoors. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 738–755. Springer, Heidelberg, April 2012. (Cited on page 2, 5, 27.)
- [Lyu16] Vadim Lyubashevsky. Digital signatures based on the hardness of ideal lattice problems in all rings. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 196–214. Springer, Heidelberg, December 2016. (Cited on page 2, 5.)
- [MV03] Daniele Micciancio and Salil P. Vadhan. Statistical zero-knowledge proofs with efficient provers: Lattice problems and more. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 282–298. Springer, Heidelberg, August 2003. (Cited on page 2.)
- [NC00] Michael A. Nielsen and Issac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000. (Cited on page 6.)
- [NIS13] NIST Special Publication 800-165 Computer Security Division 2012 Annual Report, June 2013. <http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-165.pdf> (page 39); accessed 30-January-2014. (Cited on page 2.)
- [PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In Shai Halevi and Tal Rabin, editors, *TCC 2006*, volume 3876 of *LNCS*, pages 145–166. Springer, Heidelberg, March 2006. (Cited on page 27.)
- [PS00] David Pointcheval and Jacques Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 13(3):361–396, 2000. (Cited on page 2, 27.)
- [PV05] Pascal Paillier and Damien Vergnaud. Discrete-log-based signatures may not be equivalent to discrete log. In Bimal K. Roy, editor, *ASIACRYPT 2005*, volume 3788 of *LNCS*, pages 1–20. Springer, Heidelberg, December 2005. (Cited on page 2.)
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005. (Cited on page 17.)

- [Sho04] Victor Shoup. Sequences of games: a tool for taming complexity in security proofs. Cryptology ePrint Archive, Report 2004/332, 2004. <http://eprint.iacr.org/2004/332>. (Cited on page 6.)
- [SSH11] Koichi Sakumoto, Taizo Shirai, and Harunaga Hiwatari. Public-key identification schemes based on multivariate quadratic polynomials. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 706–723. Springer, Heidelberg, August 2011. (Cited on page 5.)
- [Ste94] Jacques Stern. A new identification scheme based on syndrome decoding. In Douglas R. Stinson, editor, *CRYPTO'93*, volume 773 of *LNCS*, pages 13–21. Springer, Heidelberg, August 1994. (Cited on page 2.)
- [Unr15] Dominique Unruh. Non-interactive zero-knowledge proofs in the quantum random oracle model. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 755–784. Springer, Heidelberg, April 2015. (Cited on page 2, 3, 5.)
- [Unr17a] Dominique Unruh. Post-quantum security of fiat-shamir. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 65–95. Springer, Heidelberg, December 2017. (Cited on page 3, 4, 5.)
- [Unr17b] Dominique Unruh. Post-quantum security of fiat-shamir. Cryptology ePrint Archive, Report 2017/398, 2017. <http://eprint.iacr.org/2017/398>. (Cited on page 3, 5, 6.)
- [Zha12a] Mark Zhandry. How to construct quantum random functions. In *53rd FOCS*, pages 679–687. IEEE Computer Society Press, October 2012. (Cited on page 4, 7, 32.)
- [Zha12b] Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 758–775. Springer, Heidelberg, August 2012. (Cited on page 7.)

A Omitted Proofs

A.1 Proof of Lemma 2.1

The following lemma gives a lower bound on the number of quantum oracle queries to an oracle $g : X \rightarrow \{0, 1\}$, such that for each $x \in X$, $g(x)$ is distributed according to \mathcal{B}_λ , to find an x satisfying $g(x) = 1$. It is a reformulation of [HRS16, Theorem 1] whose proof is based on [Zha12a].

Lemma A.1 (*Generic Search Problem*) *Let $\lambda \in [0, 1]$. Then, for any (unbounded, quantum) algorithm A issuing at most Q quantum queries to $|g\rangle$, $\Pr[\text{GSP}_\lambda^A \Rightarrow 1] \leq 8 \cdot \lambda \cdot (Q + 1)^2$, where Game GSP_λ is defined in Figure 18.*

GAME GSP_λ
01 for all $x \in X$: $g(x) \leftarrow \mathcal{B}_\lambda$
02 $x \leftarrow A^{(g)}$
03 **return** $g(x)$

Figure 18: The generic quantum search game GSP_λ with Bernoulli parameter $\lambda \in [0, 1]$.

Proof of Lemma 2.1. We reduce GSPB to GSP as follows. Let $A = (A_1, A_2)$ be an adversary against GSPB_λ making Q quantum queries to oracle g_A . Let Ber_λ be a probabilistic algorithm that returns $x \leftarrow \text{Ber}_\lambda$, distributed according to the Bernoulli distribution \mathcal{B}_λ . Recall that $x := \text{Ber}_\lambda(r)$ denotes the deterministic execution of Ber_λ using explicitly given random tape r . In Figure 19 we define an adversary B against GSP_λ , also making Q quantum queries to oracle g_B . It is easy to verify that for each $x \in X$, the output distribution of g_A is $\mathcal{B}_{\lambda(x)}$. Consistency of oracle g_A is assured by deriving the randomness to sample y_A in case $y_B = 1$ (line 08) using fixed random coins $f(x)$, derived by a $2Q$ -wise independent hash

Adversary B^{g_B}	$g_A(x)$	//quantum access
01 Pick $2Q$ -wise independent hash f_{2Q}	06 $y_B := g_B(x)$	
02 $(\lambda(x))_{x \in X} \leftarrow A_1$	07 if $y_B = 0$ then $y_A := 0$	
03 if $\exists x \in X$ s.t. $\lambda(x) > \lambda$ then return 0	08 if $y_B = 1$ then $y_A := \text{Ber}_{\lambda(x)/\lambda}(f_{2Q}(x))$	
04 $x \leftarrow A_2^{g_A}$	09 return y_A	
05 return x		

Figure 19: Adversary $A = (A_1, A_2)$ in game GSPB_λ for the proof of Lemma 2.1.

function f (which looks like a perfectly random function to A). Consider the output x of the second part A_2 . Since $g_A(x) = 1$ implies $g_B(x) = 1$, we have

$$\Pr[\text{GSPB}_\lambda^A \Rightarrow 1] \leq \Pr[\text{GSP}_\lambda^B \Rightarrow 1] \leq 8(Q+1)^2\lambda.$$

This completes the proof. \square

A.2 Proof of Theorem 3.3

Proof. Let A be a quantum adversary against the UF-CMA security of SIG , issuing at most Q_H queries to $|H\rangle$, and at most Q_S queries to SIGN_1 . Consider the games given in Figure 20.

GAME G_0-G_2	SIGN(M)
01 $(pk, sk') \leftarrow \text{IGen}(\text{par})$	17 $ctr_M := ctr_M + 1$
02 $K \leftarrow \{0, 1\}^n$	18 $\mathcal{M} = \mathcal{M} \cup \{M\}$
03 $sk = (sk', K)$	19 $(W_{M,ctr_M}, c_{M,ctr_M}, Z_{M,ctr_M}) := \text{GetTrans}(M, ctr_M)$
04 $(M^*, \sigma^*) \leftarrow A^{(H\rangle), \text{SIGN}}(pk)$	20 return $\sigma_{M,ctr_M} := (W_{M,ctr_M}, Z_{M,ctr_M})$
05 Parse $\sigma^* = (W^*, Z^*)$	
06 $c^* := H(W^* \parallel M^*)$	H($W \parallel M$) //quantum access
07 if $c^* \neq H'(W^* \parallel M^*)$ then return 0	// G_2 21 for $ctr = 1$ to Q_S do
08 return $\llbracket M^* \notin \mathcal{M} \rrbracket \wedge \vee(pk, W^*, c^*, Z^*)$	22 $(W_{M,ctr}, c_{M,ctr}, Z_{M,ctr}) := \text{GetTrans}(M, ctr)$ // G_1 - G_2
	23 if $W = W_{M,ctr}$ then return $c_{M,ctr}$ // G_1 - G_2
GetTrans(M, ctr)	// G_0 24 return $H'(W \parallel M)$
09 $\kappa := 0$	GetTrans(M, ctr) // G_1 - G_2
10 while $Z_{M,ctr} = \perp$ and $\kappa \leq \kappa_m$ do	25 $\kappa := 0$
11 $\kappa := \kappa + 1$	26 while $Z_{M,ctr} = \perp$ and $\kappa \leq \kappa_m$ do
12 $(W_{M,ctr}, St) := P_1(sk; \text{RF}(0 \parallel M \parallel \kappa \parallel ctr))$	27 $\kappa := \kappa + 1$
13 $c_{M,ctr} := H(W_{M,ctr} \parallel M)$	28 $(W_{M,ctr}, c_{M,ctr}, Z_{M,ctr}) := \text{Sim}(pk; \text{RF}(M \parallel \kappa \parallel ctr))$
14 $Z_{M,ctr} := P_2(sk, W_{M,ctr}, c_{M,ctr}, St; \text{RF}(1 \parallel M \parallel \kappa \parallel ctr))$	29 if $Z_{M,ctr} = \perp$ then $(W_{M,ctr}, c_{M,ctr}) = (\perp, \perp)$
15 if $Z_{M,ctr} = \perp$ then $(W_{M,ctr}, c_{M,ctr}) = (\perp, \perp)$	30 return $(W_{M,ctr}, c_{M,ctr}, Z_{M,ctr})$
16 return $(W_{M,ctr}, c_{M,ctr}, Z_{M,ctr})$	

Figure 20: Games G_0 - G_2 for the proof of Theorem 3.3. Here RF and H' are perfect random function that cannot be accessed by A .

GAME G_0 . Note that game G_0 is the original UF-CMA game, where in lines 12 and 14 the randomness of P_1 and P_2 is derived using a perfect random function RF as follows. For each message M , the counter $1 \leq ctr_M \leq Q_S$ counts the number of classical signing queries made so far with respect to M . Counter ctr_M is then fed into RF in lines 12 and 14 to generate fresh randomness for each invocation of $\text{SIGN}(M)$. This shows

$$\Pr[G_0^A \Rightarrow 1] = \text{Adv}_{\text{SIG}}^{\text{UF-CMA}}(A).$$

GAME G_1 . This game computes the signatures on M using the naHVZK simulation algorithm Sim and patches the quantum random oracle H accordingly.

Concretely, consider the ctr_M 's classical query to $\text{SIGN}(\cdot)$ on message M and let κ_M be the smallest integer $1 \leq \kappa \leq \kappa_m$ satisfying $(W, c, Z) := \text{Sim}(pk; \text{RF}(M \parallel \kappa \parallel ctr_M))$ and $Z \neq \perp$. If no such integer

exists, then we define $\kappa_M := \perp$. It deterministically computes

$$\begin{aligned} & (W_{M,ctr_M}, c_{M,ctr_M}, Z_{M,ctr_M}) := \\ \text{GetTrans}(M, ctr_M) &= \begin{cases} \text{Sim}(pk; \text{RF}(M \parallel \kappa_M \parallel ctr_M)) & 1 \leq \kappa_M \leq \kappa_m \\ (\perp, \perp, \perp) & \kappa_M = \perp \end{cases} \end{aligned} \quad (26)$$

The signature on M is returned as

$$\sigma_{M,ctr_M} := (W_{M,ctr_M}, Z_{M,ctr_M}).$$

By the naHVZK property and the union bound, the distribution of σ_{M,ctr_M} has statistical distance at most $\kappa_m \varepsilon_{zk}$ from one computed in game G_0 . To ensure that σ_{M,ctr_M} is a valid signature on M , in line 23 the random oracle is patched such that $\text{H}(W_{M,ctr_M} \parallel M) = c_{M,ctr_M}$ holds. Concretely, a query $W \parallel M$ to quantum random oracle H with non-zero amplitude is patched with $\text{H}(W \parallel M) := c_{M,ctr_M}$ iff there exists $1 \leq ctr \leq Q_S$ such that $W = W_{M,ctr_M}$, where c_{M,ctr_M} and W_{M,ctr_M} are computed by $\text{GetTrans}(M, ctr)$. Note that the out distribution of the random oracle H in this game remains unchanged since c_{M,ctr_M} generated by the naHVZK simulator Sim is required to be uniformly distributed.

Overall, by a union bound we obtain

$$|\Pr[G_1^A \Rightarrow 1] - \Pr[G_0^A \Rightarrow 1]| \leq \kappa_m Q_S \cdot \varepsilon_{zk} .$$

GAME G_2 . This game additional defined its output as 0 in line 07 if $\text{H}(W^* \parallel M^*) \neq \text{H}'(W^* \parallel M^*)$, i.e., if $\text{H}(W^* \parallel M^*)$ was patched before in line 23. Games G_1 and G_2 can only differ if there exists $1 \leq ctr \leq Q_S$ such that $W_{M^*,ctr} = W^*$ and $M^* \notin \mathcal{M}$. Since $M^* \notin \mathcal{M}$, the random variables $W_{M^*,ctr}$ ($1 \leq ctr \leq Q_S$) were not yet revealed as part of an established signature and are completely hidden from the view of the adversary. It has α bits of min-entropy, meaning with probability at least $1 - 2^{-\alpha}$ over the keys, we have $\Pr[W_{M^*,ctr} = W^*] \leq 2^{-\alpha}$ for all $1 \leq ctr \leq Q_S$. By a union bound we obtain

$$|\Pr[G_2^A \Rightarrow 1] - \Pr[G_1^A \Rightarrow 1]| \leq Q_S 2^{-\alpha+1} .$$

To bound $|\Pr[G_2^A \Rightarrow 1]|$, consider adversary B against the UF-NMA game from Figure 21 having quantum access to random oracle H' . It perfectly simulates A 's view in game G_2 , using its own random oracle H' to simulate H and perfectly simulating the random function RF with a $2\kappa_m Q_H Q_S$ -wise independent hash function. Assume A 's forgery (M^*, σ^*) is valid in game G_2 , i.e., $M^* \notin \mathcal{M}$ and $\text{V}(pk, W^*, c^*, Z^*)$, where $c^* = \text{H}(W^* \parallel M^*)$. If $\text{H}(W^* \parallel M^*) = \text{H}'(W^* \parallel M^*)$ then (M^*, σ^*) is also a valid forgery in B 's UF-NMA game. Hence,

$$\Pr[G_2^A \Rightarrow 1] = \text{Adv}_{\text{SIG}}^{\text{UF-NMA}}(\text{B}) .$$

The running time $\text{Time}(\text{B})$ of adversary B is given by the time $\text{Time}(\text{A})$ to run A as a blackbox in game

Adversary $\text{B}^{\text{H}'}$ (pk)

01 $(M^*, \sigma^*) \leftarrow \text{A}^{\text{H}', \text{SIGN}}(pk)$

02 Parse $\sigma^* = (W^*, Z^*)$

03 $c^* := \text{H}(W^* \parallel M^*)$

04 **if** $W_{M^*} = W^*$ **then abort**

05 **if** $\llbracket M^* \notin \mathcal{M} \rrbracket \wedge \text{V}(pk, W^*, c^*, Z^*)$ **then return** (M^*, σ^*)

06 **abort**

Figure 21: Adversary B against UF-NMA security of SIG with quantum access to random oracle H' . The oracles SIGN and H simulated by B are defined as in game G_2 of Figure 20.

G_2 where in every of the Q_H oracle queries, at most $O(\kappa_m Q_S)$ computations need to be performed. The proof follows by collecting the probabilities.

STRONG UNFORGEABILITY. For sUF-CMA security we consider exactly the same games with the difference that in all games the winning condition in line 08 is changed to $\llbracket (M^*, \sigma^*) \notin \mathcal{M} \rrbracket \wedge \text{V}(pk, W^*, c^*, Z^*)$ to account for strong unforgeability, where \mathcal{M} now records all tuples (M, σ_M) of previously established messages/signature pairs.

The difference between games G_1 and G_2 is that game G_2 returns 0 in line 07 if $c^* \neq H(W^* \parallel M^*)$, i.e., if $H(W^* \parallel M^*)$ was previously patched in line 23 with $H(W^* \parallel M^*) := c_{M^*}$. Games G_1 and G_2 can only differ if there exists $1 \leq ctr^* \leq Q_S$ such that $W_{M^*, ctr^*} = W^*$, $M^* \notin \mathcal{M}$, and $V(pk, W^*, c^*, Z^*) = 1$. (In that case G_2 returns 0 and G_1 returns 1.)

We distinguish two cases. If $(M^*, \cdot) \notin \mathcal{M}$ then we are in the situation that the adversary did not query a signature on M^* and we can use the same argument as in standard unforgeability to argue $|\Pr[G_2^A \Rightarrow 1] - \Pr[G_1^A \Rightarrow 1]| \leq Q_S \cdot 2^{-\alpha+1}$. It leaves to handle the case $(M^*, \cdot) \in \mathcal{M}$, i.e., the adversary obtained a signatures $\sigma_{M^*} = (W_{M^*}, Z_{M^*})$ on message M^* and submits a correct forgery $\sigma^* = (W^*, Z^*)$ satisfying $W^* = W_{M^*, ctr^*}$ and $Z^* \neq Z_{M^*, ctr^*}$ for some $1 \leq ctr^* \leq Q_S$. The problem of finding values $(W^*, c^*, Z_{M^*, ctr^*}, Z^*)$ with two accepting transcripts (W^*, c^*, Z^*) and $(W^*, c^*, Z_{M^*, ctr^*})$ is exactly bounded by the advantage of an adversary C against the CUR experiment, i.e., $|\Pr[G_2^A \Rightarrow 1] - \Pr[G_1^A \Rightarrow 1]| \leq \text{Adv}_{\text{ID}}^{\text{CUR}}(\text{C})$.

In combination this proves

$$|\Pr[G_2^A \Rightarrow 1] - \Pr[G_1^A \Rightarrow 1]| \leq Q_S \cdot 2^{-\alpha+1} + \text{Adv}_{\text{ID}}^{\text{CUR}}(\text{C}).$$

Finally, a straightforward modification of adversary B against UF-NMA security to account for the strong unforgeability check proves

$$\Pr[G_2^A \Rightarrow 1] = \text{Adv}_{\text{SIG}}^{\text{UF-NMA}}(\text{B})$$

and completes proof of sUF-CMA security. \square

B A Simple ID / Signature Scheme and a Security Proof Sketch

The identification scheme which we will be converting into a signature scheme is presented in Figure 22. We do not set concrete parameters, since we will only provide a high-level sketch of the proof of its security, but the reader can look at the Dilithium-QROM column of Table 1. Other than the sizes of the keys and signatures (which will be larger for our simple scheme), the other parameters remain virtually the same.

I Gen(par)	P ₁ (sk)
01 $\rho \leftarrow \{0, 1\}^{256}$	08 $\mathbf{A} \leftarrow R_q^{k \times k} := \text{Sam}(\rho)$
02 $\mathbf{A} \leftarrow R_q^{k \times k} := \text{Sam}(\rho)$	09 $\mathbf{y}_1, \mathbf{y}_2 \leftarrow S_\gamma^k$
03 $\mathbf{s}_1, \mathbf{s}_2 \leftarrow S_\eta^k$	10 $\mathbf{w} := \mathbf{A}\mathbf{y}_1 + \mathbf{y}_2$
04 $\mathbf{t} := \mathbf{A}\mathbf{s}_1 + \mathbf{s}_2$	11 return ($W = \mathbf{w}, St = (\mathbf{w}, \mathbf{y}_1, \mathbf{y}_2)$)
05 $pk = (\rho, \mathbf{t})$	
06 $sk = (\rho, \mathbf{s}_1, \mathbf{s}_2)$	P ₂ (sk, $W = \mathbf{w}, c, St = (\mathbf{w}, \mathbf{y}_1, \mathbf{y}_2)$)
07 return (pk, sk)	12 $\mathbf{z}_1 := \mathbf{y}_1 + c\mathbf{s}_1, \mathbf{z}_2 := \mathbf{y}_2 + c\mathbf{s}_2$
	13 if $\ \mathbf{z}_1\ _\infty \geq \gamma - \beta$ or $\ \mathbf{z}_2\ _\infty \geq \gamma - \beta$
	14 then $(\mathbf{z}_1, \mathbf{z}_2) := \perp$
	15 return $Z = (\mathbf{z}_1, \mathbf{z}_2)$
V (pk, $W = \mathbf{w}, c, Z = (\mathbf{z}_1, \mathbf{z}_2)$)	
16 return $\llbracket \mathbf{w} = \mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 - c\mathbf{t} \rrbracket$ and $\llbracket \ \mathbf{z}_1\ _\infty < \gamma - \beta \rrbracket$ and $\llbracket \ \mathbf{z}_2\ _\infty < \gamma - \beta \rrbracket$	

Figure 22: A simple, unoptimized version of our ID scheme – a concrete instantiation based on the hardness of the MLWE problem of the commitment-recoverable (Definition 2.4) canonical identification scheme in Figure 4.

It is not hard to verify that the protocol is correct in the case that the Prover doesn't send \perp . One can also compute the probability that \perp is sent using Equation (12) in Lemma 4.4 for \mathbf{z}_1 and \mathbf{z}_2 .

We will now sketch the proofs of the properties that are needed for the protocol in 22 to be converted to a quantum-secure digital signature scheme. We first need to prove the identification scheme is perfectly naHVZK, i.e., the distribution of the output of the Trans algorithm (Figure 23, left) that uses the secret key as input is exactly that of the Sim algorithm (Figure 23, right) that uses only the public key as input. This involves showing that the distribution of $\mathbf{z}_1, \mathbf{z}_2$ is independent of the exact sk . This proof is fairly standard, and we give it for an analogous statement for our optimized scheme in Lemma 4.3. The proof

of what we need for the non-optimized scheme comes from applying Equation (9) twice. This will prove that the ε_{zk} in the statement of Theorem 3.1 is 0.

<p>Algorithm Trans(sk):</p> <pre> 01 $\mathbf{A} \leftarrow R_q^{k \times k} := \text{Sam}(\rho)$ 02 $\mathbf{y}_1, \mathbf{y}_2 \leftarrow S_{\gamma-1}^k$ 03 $\mathbf{w} := \mathbf{A}\mathbf{y}_1 + \mathbf{y}_2$ 04 $c \leftarrow \text{ChSet}$ 05 $\mathbf{z}_1 := \mathbf{y}_1 + c\mathbf{s}_1, \mathbf{z}_2 := \mathbf{y}_2 + c\mathbf{s}_2$ 06 if $\ \mathbf{z}_1\ _\infty \geq \gamma - \beta$ or $\ \mathbf{z}_2\ _\infty \geq \gamma - \beta$, then return \perp 07 return $(c, (\mathbf{z}_1, \mathbf{z}_2))$ </pre>	<p>Algorithm Sim(pk):</p> <pre> 08 $\mathbf{A} \leftarrow R_q^{k \times k} := \text{Sam}(\rho)$ 09 with probability $1 - \left(1 - \frac{ S_{\gamma-\beta-1}^k }{ S_{\gamma-1}^k }\right)^2$, return \perp 10 $\mathbf{z}_1, \mathbf{z}_2 \leftarrow S_{\gamma-\beta-1}^k$ 11 $c \leftarrow \text{ChSet}$ 12 return $(c, (\mathbf{z}_1, \mathbf{z}_2))$ </pre>
---	--

Figure 23: Left: a real transcript of our simple scheme output by the transcript algorithm $\text{Trans}(sk)$; Right: a simulated transcript output by the $\text{Sim}(pk)$ algorithm.

We next have to show lossiness. In particular, we need to show that if the public key generated as in Figure 24, then no (unbounded) Adversary C can win the LOSSY-IMP game from Figure 6, which is adapted to our scheme in Figure 25, with noticeable probability.

<p>LossyGen(par)</p> <pre> 01 $\rho \leftarrow \{0, 1\}^{256}, \mathbf{A} \leftarrow R_q^{k \times k} := \text{Sam}(\rho)$ 02 $\mathbf{t} \leftarrow R_q^k$ 03 return $pk = (\rho, \mathbf{t})$ </pre>

Figure 24: The lossy instance generator LossyGen for the simple scheme.

<p>GAME LOSSY-IMP:</p> <pre> 01 $pk_{\text{is}} := (\rho, \mathbf{t}) \leftarrow \text{LossyGen}(\text{par})$ 02 $(\mathbf{w}, St) \leftarrow C(pk_{\text{is}})$ 03 $c \leftarrow \text{ChSet}$ 04 $(\mathbf{z}_1, \mathbf{z}_2) \leftarrow C(St, c)$ 05 return $[\mathbf{w} = \mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 - c\mathbf{t}]$ and $[\ \mathbf{z}_1\ _\infty < \gamma - \beta]$ and $[\ \mathbf{z}_2\ _\infty < \gamma - \beta]$ </pre>
--

Figure 25: The lossy impersonation game LOSSY-IMP for the simple scheme.

The intuition for proving this follows the framework from [KW03, AFLT12]. We want to set the parameters such that for all but a negligible fraction of (\mathbf{A}, \mathbf{t}) , for all \mathbf{w} , there is only one c for which we can find $\mathbf{z}_1, \mathbf{z}_2$ that satisfy the verification equation. The proof uses a counting argument. If for some \mathbf{A}, \mathbf{t} , there is a \mathbf{w} such that $\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 - c\mathbf{t} = \mathbf{w} = \mathbf{A}\mathbf{z}'_1 + \mathbf{z}'_2 - c'\mathbf{t}$, then this implies that

$$\mathbf{A}(\mathbf{z}_1 - \mathbf{z}'_1) + (\mathbf{z}_2 - \mathbf{z}'_2) = (c - c')\mathbf{t},$$

where $\mathbf{z}_i, \mathbf{z}'_i$ and c, c' have small coefficients. For convenience, this can be rewritten as

$$\mathbf{A}\mathbf{r}_1 + \mathbf{r}_2 = d\mathbf{t}, \tag{27}$$

where $d \neq 0$. If the parameters are set such that all polynomials with small coefficients are invertible (and so \mathbf{r}_1 is invertible), then one can show that for each fixed $\mathbf{r}_1, \mathbf{r}_2, d$, the probability over the choice of \mathbf{A}, \mathbf{t} that Equation (27) is satisfied is at most q^{-kn} . Thus if the total number of possible $\mathbf{r}_1, \mathbf{r}_2, d$ is significantly less than that, then using the union bound concludes the proof. This proves that the ε_{is} term in Theorem 3.1 is very close to $1/|\text{ChSet}|$.

Lower bounding the min entropy of \mathbf{w} (i.e. showing that the α term in Theorem 3.1 is negligible) and showing that the scheme has the Computational Unique Response (CUR) property is done using

virtually the same counting argument as above. Therefore any adversary who can provide a forgery to the signature scheme in Figure 26 can also be used to produce a forgery in the identification scheme with virtually the same advantage and time. Forgery of the identification scheme is, in turn, as hard as solving the MLWE problem.

<pre> Sign((sk, K), M) 01 $\kappa := 0$ 02 $\mathbf{A} \leftarrow R_q^{k \times k} := \text{Sam}(\rho)$ 03 while $(\mathbf{z}_1, \mathbf{z}_2) = \perp$ and $\kappa \leq 200/(1 - \delta)$ do 04 $\kappa := \kappa + 1$ 05 $\mathbf{y}_1, \mathbf{y}_2 \leftarrow S_{\gamma-1}^k := \text{Sam}(K \parallel M \parallel \kappa)$ 06 $\mathbf{w} := \mathbf{A}\mathbf{y}_1 + \mathbf{y}_2$ 07 $c := \text{H}(\mathbf{w} \parallel M)$ 08 $\mathbf{z}_1 := \mathbf{y}_1 + c\mathbf{s}_1, \mathbf{z}_2 := \mathbf{y}_2 + c\mathbf{s}_2$ 09 if $\ \mathbf{z}_1\ _\infty \geq \gamma - \beta$ or $\ \mathbf{z}_2\ _\infty \geq \gamma - \beta$ then $(\mathbf{z}_1, \mathbf{z}_2) := \perp$ 10 return $\sigma = (\mathbf{z}_1, \mathbf{z}_2, c)$ Ver(pk, M, $\sigma = (\mathbf{z}, \mathbf{h}, c)$) 11 $\mathbf{A} \leftarrow R_q^{k \times k} := \text{Sam}(\rho)$ 12 $\mathbf{w} := \mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 - c\mathbf{t}$ 13 return $\llbracket \ \mathbf{z}_1\ _\infty < \gamma - \beta \rrbracket$ and $\llbracket \ \mathbf{z}_2\ _\infty < \gamma - \beta \rrbracket$ and $\llbracket c = \text{H}(\mathbf{w}_1 \parallel M) \rrbracket$ </pre>

Figure 26: The simple signature scheme. The key generation algorithm is IGen from Figure 22, where the secret key also contains a random key K for the pseudorandom function $\text{Sam}(K \parallel \cdot)$.

C Lower-bounding the min-entropy α of Dilithium

In Dilithium-QROM, we were able to lower-bound α , the min-entropy of the identification scheme, by using Lemma 4.7. In Dilithium, the conditions of the Lemma are no longer satisfied, and below we give an alternate proof for bounding α .

Lemma C.1

$$\Pr_{\mathbf{A} \leftarrow R_q^{k \times \ell}} \left[\forall \mathbf{w}_1 : \Pr_{\mathbf{y} \leftarrow S_{\gamma'-1}^\ell} [\text{HighBits}_q(\mathbf{A}\mathbf{y}, 2\gamma) = \mathbf{w}_1] \leq \left(\frac{2\gamma + 1}{2\gamma' - 1} \right)^n \right] > 1 - (n/q)^{k\ell}.$$

Proof. The probability that a random polynomial $a \leftarrow R_q$ is invertible in $R_q = \mathbb{Z}_q[X]/(X^n + 1)$ when the polynomial $X^n + 1$ splits into n linear factors is $(1 - 1/q)^n > 1 - n/q$. Thus the probability that at least one of $k\ell$ polynomials in $\mathbf{A} \leftarrow R_q^{k \times \ell}$ is invertible is greater than $1 - (n/q)^{k\ell}$. We will now prove that for all \mathbf{A} that contain at least one invertible polynomial, we will have that for all \mathbf{w}_1 ,

$$\Pr_{\mathbf{y} \leftarrow S_{\gamma'-1}^\ell} [\text{HighBits}_q(\mathbf{A}\mathbf{y}, 2\gamma) = \mathbf{w}_1] \leq \left(\frac{2\gamma + 1}{2\gamma' - 1} \right)^n,$$

which will prove the Lemma. Let us only consider the row of \mathbf{A} which contains the invertible polynomial. Call the elements in this row $[a_1, \dots, a_\ell]$ and without loss of generality assume that a_1 is invertible. We will want to prove that for all w_1 ,

$$\Pr_{\mathbf{y} \leftarrow S_{\gamma'-1}^\ell} \left[\text{HighBits}_q \left(\sum a_i y_i, 2\gamma \right) = w_1 \right] \leq \left(\frac{2\gamma + 1}{2\gamma' - 1} \right)^n.$$

Define T to be the set containing all the elements w such that $\text{HighBits}(w, 2\gamma) = w_1$. By the definition of the Decompose_q routine in Figure 12, the size of T is at most $(2\gamma + 1)^n$. We can then rewrite the above probability as

$$\Pr_{\mathbf{y} \leftarrow S_{\gamma'-1}^\ell} \left[\sum a_i y_i \in T \right] = \Pr_{y_1 \leftarrow S_{\gamma'-1}^\ell} \left[y_1 \in a_1^{-1} \left(T - \sum a_i y_i \right) \right] \leq \left(\frac{2\gamma + 1}{2\gamma' - 1} \right)^n,$$

where the last inequality follows due to the fact that the size of the set

$$a_1^{-1} \left(T - \sum a_i y_i \right)$$

is the same as that of T , which is at most $(2\gamma + 1)^n$; and the size of the set $S_{\gamma'-1}$ is exactly $(2\gamma' - 1)^n$. \square

For the values in Table 1, we have that $\left(\frac{2\gamma+1}{2\gamma'-1}\right)^n < 2^{-255}$ and $(n/q)^{k\ell} < 2^{-299}$. Thus, by Definition 2.6, the min-entropy of the Dilithium scheme for the two sets of parameters is greater than 255.

We would like to point out that the real min-entropy should be a lot higher since the HighBits_q function maps onto a set of size larger than 2^{5000} and is heuristically close to uniform over this set. To get a formal proof would be significantly more involved than the proof above which took advantage of the fact that $\gamma' = 2\gamma$, and gave us a sufficiently high min-entropy bound for practical purposes.