

Secure Signatures and Chosen Ciphertext Security in a Post-Quantum World

Dan Boneh Mark Zhandry

Stanford University
{dabo,zhandry}@cs.stanford.edu

Abstract

We initiate the study of *quantum*-secure digital signatures and *quantum* chosen ciphertext security. In the case of signatures, we enhance the standard chosen message query model by allowing the adversary to issue *quantum* chosen message queries: given a superposition of messages, the adversary receives a superposition of signatures on those messages. Similarly, for encryption, we allow the adversary to issue *quantum* chosen ciphertext queries: given a superposition of ciphertexts, the adversary receives a superposition of their decryptions. These adversaries model a natural post-quantum environment where end-users sign messages and decrypt ciphertexts on a personal quantum computer.

We construct classical systems that remain secure when exposed to such quantum queries. For signatures we construct two compilers that convert classically secure signatures into signatures secure in the quantum setting and apply these compilers to existing post-quantum signatures. We also show that standard constructions such as Lamport one-time signatures and Merkle signatures remain secure under quantum chosen message attacks, thus giving signatures whose quantum security is based on generic assumptions. For encryption, we define security under quantum chosen ciphertext attacks and present both public-key and symmetric-key constructions.

Keywords: Quantum computing, signatures, encryption, post-quantum security

1 Introduction

Recent progress in building quantum computers [IBM12] gives hope for their eventual feasibility. Consequently, there is a growing need for post-quantum cryptosystems, namely classical systems that remain secure against quantum computers. Post quantum cryptography generally studies the settings where the adversary is armed with a quantum computer, but users only have classical machines. In this paper we go a step further and study the eventuality where end-user machines are quantum. In these settings an attacker may interact with honest parties using quantum queries, as discussed below, potentially giving the attacker more power. The challenge is to construct cryptosystems that remain secure when exposed to such quantum queries. We emphasize that all the systems we consider are classical and can be easily implemented on a classical computer. Our goal is to construct classical systems that remain secure even when implemented on a quantum computer thereby potentially giving the attacker the ability to issue quantum queries.

Along these lines, Zhandry [Zha12b] showed how to construct pseudorandom functions (PRFs) that remain secure even when the adversary is allowed to issue *quantum* queries to the PRF. A

quantum query is a superposition of inputs $\sum_x \psi_x |x\rangle$ of the attacker’s choice. The response is a superposition $\sum_x \psi_x |x, F(k, x)\rangle$ where $F(k, x)$ is the value of the PRF at a point x under key k . Zhandry showed that certain PRFs are secure even under such a powerful query model. More recently, Boneh and Zhandry [BZ13] showed how to construct message authentication codes (MACs) that remain secure even when the attacker is allowed to issue *quantum* chosen message queries. That is, for a superposition of messages $\sum_m \psi_m |m\rangle$ of the attacker’s choice, the attacker is given $\sum_m \psi_m |m, S(k, m)\rangle$ where $S(k, m)$ is the tag on message m using key k . They showed that some classically secure MACs become insecure under quantum chosen message queries and they constructed several quantum-secure MAC families.

Our contributions. In this paper, we construct the first quantum-secure signatures and quantum-secure chosen ciphertext encryption systems.

We begin by defining security for digital signatures under a *quantum* chosen message attack. A quantum chosen message query [BZ13] gives the attacker the signatures on all messages in a quantum superposition. In more detail, a quantum chosen message query is the transformation

$$\sum_m \psi_m |m\rangle \quad \longrightarrow \quad \sum_m \psi_m |m, S(\text{sk}, m)\rangle$$

where $S(\text{sk}, x)$ is the signature on x using signing key sk . The attacker can sample the response to such a query and obtain one valid message-signature pair. After q such queries, it can obtain q valid message-signature pairs. We say that a signature scheme is existentially unforgeable under a *quantum* chosen message attack if, after q quantum chosen message queries, the attacker cannot produce $q + 1$ valid message-signature pairs.

Next, we present several compilers that convert a signature scheme that is secure under *classical* queries into one secure under *quantum* queries. In particular, we give the following constructions:

- Using a chameleon hash [KR00], we show that any existentially unforgeable signature secure under a *classical* random message attack can be made existentially unforgeable under a *quantum* chosen message attack. We apply this conversion to several existing signature schemes, giving constructions whose quantum security is based on the quantum hardness of lattice problems.
- We show that any *universally* unforgeable signature under a *classical* random message attack can be made *existentially* unforgeable under a *quantum* chosen message attack in the random oracle model. For example, this conversion applies to a randomized variant of GPV signatures [GPV08], proving security of the scheme even under a *quantum* chosen message attack. We also separately show that the basic deterministic GPV scheme is secure in this setting.
- Finally, we prove that classical constructions such as Lamport one-time signatures and Merkle signatures are existentially unforgeable under a *quantum* chosen message attack. The proofs together show how to build quantum-secure signatures from any collision resistant hash function. We leave open the problem of basing security on one-way functions. We also note that the version of Lamport signatures that we prove secure are non-optimized, and can potentially be made more efficient using standard combinatorial techniques. Unfortunately, we cannot prove that these simple optimizations are quantum-secure and leave that as an interesting open problem.

Turning to encryption, we first explain how to adapt the chosen ciphertext security game to the quantum setting. In the classical game, the attacker is given classical access to a decryption oracle used to answer chosen ciphertext queries and to an encryption oracle used to create the challenge ciphertext. In the quantum settings, the decryption oracle accepts a superposition of ciphertexts and returns a superposition of their decryptions:

$$\sum_m \psi_c |c\rangle \quad \longrightarrow \quad \sum_c \psi_c |c, D(\text{sk}, c)\rangle .$$

One might also try to allow quantum access to the encryption oracle — however, as we show, that avenue results in unsatisfiable definitions. We therefore restrict the encryption oracle to be classical.

Armed with this definition of security, we construct quantum-secure chosen ciphertext systems in both the public-key and symmetric-key settings:

- Our symmetric-key construction is built from any secure PRF, and follows the encrypt-then-MAC paradigm. The classical proof that encrypt-then-MAC is secure for generic encryption and generic MAC schemes does not carry over to the quantum setting, but we are able to prove security for our specific construction.
- We show that public-key quantum chosen ciphertext security can be obtained from any identity-based encryption scheme that is selectively secure under a quantum chosen identity attack. Such an identity-based encryption scheme can, in turn, be built from lattice assumptions.

Motivation. Allowing the adversary to issue quantum queries is a natural and conservative security model and is therefore an interesting one to study. Constructing signature and encryption schemes that remain secure in these models gives confidence in the event that end-user computing devices eventually become quantum. Nevertheless, one might imagine that in a future where computers are quantum, the last step in a signature or decryption procedure is to sample the final quantum state so that the results are always classical. The quantum query model ensures that even if the attacker can bypass this last “classicalization” step, the cryptosystem remains secure.

Other related work. Several recent works study the security of cryptographic primitives when the adversary can issue quantum queries. Boneh et al. [BDF⁺11] and Zhandry [Zha12a] prove the classical security of signatures, encryption, and identity-based encryption schemes in the *quantum* random oracle model, where the adversary can query the random oracle on superpositions of inputs. In these papers, the interaction with the challenger is classical. These results show that many, but not all, random oracle constructions remain secure in the quantum random oracle model. The quantum random oracle model has also been used to prove security of Merkle’s Puzzles in the quantum settings [BS08, BHK⁺11]. Meanwhile, Damård et al. [DFNS11] examine secret sharing and multiparty computation in a model where an adversary may corrupt a superposition of subsets of players, and build zero knowledge protocols that are secure, even when a dishonest verifier can issue challenges on superpositions.

Some progress toward identifying sufficient conditions under which classical protocols are also quantum immune has been made by Unruh [Unr10] and Hallgren et al. [HSS11]. Unruh shows that any scheme that is statistically secure in Cannetti’s universal composition (UC) framework [Can01] against classical adversaries is also statistically secure against quantum adversaries. Hallgren et al. show that for many schemes, this is also true in the computational setting. These results, however, do not apply to cryptographic primitives such as signatures and encryption.

2 Preliminaries: Background and Techniques

We will let $[n]$ denote the set $\{1, \dots, n\}$. Functions will be denoted by capital letters (such as F), and sets by capital script letters (such as \mathcal{X}). Given a function $F : \mathcal{X} \rightarrow \mathcal{Y}$ and a subset $\mathcal{S} \subseteq \mathcal{X}$, the restriction of F to \mathcal{S} is the function $F_{\mathcal{S}} : \mathcal{S} \rightarrow \mathcal{Y}$ where $F_{\mathcal{S}}(x) = F(x)$ for all $x \in \mathcal{S}$. A distribution D on F induces a distribution $D_{\mathcal{S}}$ on $F_{\mathcal{S}}$. We say that D is k -wise independent if each of the distributions $D_{\mathcal{S}}$ are truly random distributions on functions from \mathcal{S} to \mathcal{Y} , for all sets \mathcal{S} of size at most k . A set \mathcal{F} of functions from \mathcal{X} to \mathcal{Y} is k -wise independent if the uniform distribution on \mathcal{F} is k -wise independent.

2.1 Quantum Computation

We give a short introduction to quantum computation. A quantum system A is a complex Hilbert space \mathcal{H} together with an inner product $\langle \cdot | \cdot \rangle$. The state of a quantum system is given by a vector $|\psi\rangle$ of unit norm ($\langle \psi | \psi \rangle = 1$). Given quantum systems \mathcal{H}_1 and \mathcal{H}_2 , the joint quantum system is given by the tensor product $\mathcal{H}_1 \otimes \mathcal{H}_2$. Given $|\psi_1\rangle \in \mathcal{H}_1$ and $|\psi_2\rangle \in \mathcal{H}_2$, the product state is given by $|\psi_1\rangle|\psi_2\rangle \in \mathcal{H}_1 \otimes \mathcal{H}_2$. Given a quantum state $|\psi\rangle$ and an orthonormal basis $B = \{|0\rangle, \dots, |d-1\rangle\}$ for \mathcal{H} , a measurement of $|\psi\rangle$ in the basis B results in the vector $|b_i\rangle$ with probability $|\langle b_i | \psi \rangle|^2$.

A unitary transformation over a d -dimensional Hilbert space \mathcal{H} is a $d \times d$ matrix \mathbf{U} such that $\mathbf{U}\mathbf{U}^\dagger = \mathbf{I}_d$, where \mathbf{U}^\dagger represents the conjugate transpose. A quantum algorithm operates on a product space $\mathcal{H}_{in} \otimes \mathcal{H}_{out} \otimes \mathcal{H}_{work}$ and consists of n unitary transformations $\mathbf{U}_1, \dots, \mathbf{U}_n$ in this space. \mathcal{H}_{in} represents the input to the algorithm, \mathcal{H}_{out} the output, and \mathcal{H}_{work} the work space. A classical input x to the quantum algorithm is converted to the quantum state $|x, 0, 0\rangle$. Then, the unitary transformations are applied one-by-one, resulting in the final state

$$|\psi_x\rangle = \mathbf{U}_n \dots \mathbf{U}_1 |x, 0, 0\rangle .$$

The final state is then measured, obtaining (a, b, c) with probability $|\langle a, b, c | \psi_x \rangle|^2$. The output of the algorithm is b . We say that a quantum algorithm is efficient if each of the unitary matrices \mathbf{U}_i come from some fixed basis set, and n , the number of unitary matrices, is polynomial in the size of the input.

Quantum-accessible Oracles. We will implement an oracle $O : \mathcal{X} \rightarrow \mathcal{Y}$ by a unitary transformation \mathbf{O} where

$$\mathbf{O}|x, y, z\rangle = |x, y + O(x), z\rangle$$

where $+$: $\mathcal{X} \times \mathcal{X} \rightarrow \mathcal{X}$ is some group operation on \mathcal{X} . Suppose we have a quantum algorithm that makes quantum queries to oracles O_1, \dots, O_q . Let $|\psi_0\rangle$ be the state of the algorithm before any queries, and let $\mathbf{U}_1, \dots, \mathbf{U}_q$ be the unitary transformations applied between queries. The final state of the algorithm will be

$$\mathbf{U}_q \mathbf{O}_q \dots \mathbf{U}_1 \mathbf{O}_1 |\psi_0\rangle$$

We can also have an algorithm make classical queries to O_i . In this case, the input to the oracle is measured before applying the transformation \mathbf{O}_i .

Fix an oracle $O : \mathcal{X} \rightarrow \mathcal{Y}$. Let $O^{(q)} : \mathcal{X}^q \rightarrow \mathcal{Y}^q$ be the oracle that maps \mathbf{x} into $O(\mathbf{y}) = (O(x_1), O(x_2), \dots, O(x_q))$. Observe that any quantum query to $O^{(q)}$ can be implemented using q quantum queries to O , where the unitary transformations between queries just permute the registers. We say that an algorithm that makes a single query to $O^{(q)}$ makes q non-adaptive queries to O .

Tools. Next we state several lemmas and definitions that we will use throughout the paper. Some have been proved in other works, and the rest are proved in Appendix B. The first concerns partial measurements, and will be used extensively throughout the paper:

Lemma 2.1. *Let A be a quantum algorithm, and let $\Pr[x]$ be the probability that A outputs x . Let A' be another quantum algorithm obtained from A by pausing A at an arbitrary stage of execution, performing a partial measurement that obtains one of k outcomes, and then resuming A . Let $\Pr'[x]$ be the probability A' outputs x . Then $\Pr'[x] \geq \Pr[x]/k$.*

This lemma means, for example, that if you measure just one qubit, the probability of a particular output drops by at most a factor of two. We also make use of the following lemma, proved by Zhandry [Zha12a], which allows us to simulate random oracle efficiently using k -wise independent functions:

Lemma 2.2 ([Zha12a]). *Let H be an oracle drawn from a $2q$ -wise independent distribution. Then H looks like a random function to any quantum algorithm making at most q quantum queries to H .*

The next definition and lemma are given by Zhandry [Zha12b] and allow for the efficient simulation of an exponentially-large list of samples, given only a polynomial number of samples:

Definition 2.3 (Small-range distributions [Zha12b]). *Fix sets \mathcal{X} and \mathcal{Y} and a distribution D on \mathcal{Y} . Fix an integer r . Let $\mathbf{y} = (y_1, \dots, y_r)$ be a list of r samples from D and let P be a random function from \mathcal{X} to $[r]$. Then the distribution of functions $H : \mathcal{X} \rightarrow \mathcal{Y}$ defined by $H(x) = y_{P(x)}$ is called a small-range distribution with r samples of D .*

Lemma 2.4 ([Zha12b]). *There is a universal constant C_0 such that, for any sets \mathcal{X} and \mathcal{Y} , distribution D on \mathcal{Y} , any integer ℓ , and any quantum algorithm A making q queries to an oracle $H : \mathcal{X} \rightarrow \mathcal{Y}$, the following two cases are indistinguishable, except with probability less than $C_0 q^3 / \ell$:*

- $H(x) = y_x$ where \mathbf{y} is a list of samples of D of size $|\mathcal{X}|$.
- H is drawn from the small-range distribution with ℓ samples of D .

We use Lemma 2.4 to prove the following corollary:

Lemma 2.5. *Let \mathcal{X} and \mathcal{Y} be sets, and for each $x \in \mathcal{X}$, let D_x and D'_x be distributions on \mathcal{Y} such that $|D_x - D'_x| \leq \epsilon$ for some fixed ϵ . Let $O : \mathcal{X} \rightarrow \mathcal{Y}$ be a function where, for each x , $O(x)$ is drawn from D_x , and let $O'(x)$ be a function where, for each x , $O'(x)$ is drawn from D'_x . Then any quantum algorithm making at most q queries to either O or O' cannot distinguish the two, except with probability at most $\sqrt{8C_0 q^3 \epsilon}$.*

Zhandry [Zha12b] proves this corollary for the special case where all of the D_x distributions are the same and all of the D'_x distributions are the same. Lastly, we need the following lemma:

Lemma 2.6. *Fix sets \mathcal{X} and \mathcal{Y} , and a distribution D on \mathcal{Y} . Let H be a function from \mathcal{X} to \mathcal{Y} where each output is drawn independently according to D . Let F be a function from \mathcal{Y} to some set \mathcal{Z} . Then any quantum algorithm making q quantum queries to H and an arbitrary number of queries to $F \circ H$ can only produce $q + 1$ input/output pairs of H with probability at most $(q + 1) / \left\lfloor 2^{H_\infty} \right\rfloor$, where H_∞ be the minimum over all $z \in \mathcal{Z}$ of the min-entropy of the distribution D on y conditioned on $F(y) = z$.*

A special case of this theorem is when F is a constant function and D is uniform. In this case, Lemma 2.6 reduces to the following result of Boneh and Zhandry [BZ13]: any quantum algorithm making q queries to a random oracle H from \mathcal{X} to \mathcal{Y} can output $q + 1$ input/output pairs of H with probability at most $(q + 1)/|\mathcal{Y}|$. We prove Lemma 2.6 by reducing the general case to this special case with $|\mathcal{Y}| = \lfloor 2^{H_\infty} \rfloor$.

3 Signatures

We first define signatures in the quantum setting and what it means for them to be secure.

Definition 3.1. *A signature scheme \mathcal{S} is a tuple of algorithms $(G, \text{Sign}, \text{Ver})$ where*

- $G(\lambda)$ generates a private/public key pair (sk, pk) .
- $\text{Sign}(\text{sk}, m)$ is a potentially stateful algorithm that generates a signature σ . If Sign is stateful, we require that the state does not depend in any way on the messages that have been signed.
- $\text{Ver}(\text{pk}, m, \sigma)$ either accepts or rejects. We require that valid signatures always accept, that is $\text{Ver}(\text{pk}, m, \text{Sign}(\text{sk}, m))$ accepts.

The requirement that the state of the signing algorithm does not depend on previously signed messages is new relative to classical definitions for stateful signatures. This is necessary for the state of the signing algorithm to remain classical when signing a superposition of messages. Our definition of security for signatures will be similar to the security notion for message authentication codes defined by Boneh and Zhandry [BZ13]:

Definition 3.2 (Quantum Security). *A signature scheme $\mathcal{S} = (G, \text{Sign}, \text{Ver})$ is strongly existentially unforgeable under a quantum chosen-message attack (EUF-qCMA secure) if, for any efficient quantum algorithm A and any polynomial q , A 's probability of success in the following game is negligible in λ :*

Key Gen *The challenger runs $(\text{sk}, \text{pk}) \leftarrow G(\lambda)$, and gives pk to A .*

Signing Queries *The adversary makes a polynomial q chosen message queries. For each query, the challenger chooses randomness r , and responds by signing each message in the query using r as randomness:*

$$\sum_{m,t} \psi_{m,t} |m, t\rangle \quad \longrightarrow \quad \sum_{m,t} \psi_{m,t} |m, t \oplus \text{Sign}(\text{sk}, m; r)\rangle$$

Forgeries *The adversary is required to produce $q + 1$ message/signature pairs. The challenger then checks that all the signatures are valid, and that all message/signature pairs are distinct. If so, the challenger reports that the adversary wins.*

We also give some weaker security definitions. The first is for random message security:

Definition 3.3. *\mathcal{S} is existentially unforgeable under a random message attack (EUF-RMA secure) if the adversary is not allowed any signing queries, but instead receives q message/signature pairs for random messages at the beginning of the game.*

We can weaken the security definition even further, to get universal unforgeability:

Definition 3.4. \mathcal{S} is universally unforgeable under a random message attack (UUF-RMA secure) if, in addition to receiving q message/signature pairs for random messages, the adversary also receives n random messages, and all of the messages forged must be among the $q + n$ messages received.

All of the above security definition also have weak variants, where in addition to requiring the message/signature forgery pairs to be distinct, we require all the messages themselves to be distinct. Finally, all of the above security definitions also have k -time variants for any constant k , where the value of q is bounded to at most k . When the distinction is required, we refer to the standard variants as many-time security.

3.1 Quantum-Secure Signatures from Classically-Secure Signatures

In this section, we show a general transformation from classically secure signatures to quantum secure signatures. The building blocks for our construction are chameleon hash functions and signatures that are secure against a classical random message attack. First, we will define a chameleon hash function. The definition we use is slightly different from the original definition from Krawczyk and Rabin [KR00], but is satisfied by the known lattice constructions:

Definition 3.5. A chameleon hash function \mathcal{H} is a tuple of algorithms $(G, H, \text{Inv}, \text{Sample})$ where:

- $G(\lambda)$ generates a secret/public key pair (sk, pk) .
- $H(\text{pk}, m, r)$ maps messages to some space \mathcal{Y}
- $\text{Sample}(\lambda)$ samples r from some distribution such that, for every pk and m , $H(\text{pk}, m, r)$ is uniformly distributed.
- $\text{Inv}(\text{sk}, h, m)$ produces an r such that $H(\text{pk}, m, r) = h$, and r is distributed negligibly-close to $\text{Sample}(\lambda)$ conditioned on $H(\text{pk}, m, r) = h$

We say that a chameleon hash function is collision resistant if no efficient quantum algorithm, given only pk , can find collisions in $H(\text{pk}, \cdot, \cdot)$. Cash et al.[CHKP10] build a simple lattice-based chameleon hash function, and prove that it is collision resistant, provided that the *Shortest Integer Solution* problem (SIS) is hard for an appropriate choice of parameters. In what follows, for any randomized algorithm A , we let $A(x; r)$ denote running A on input x with randomness r .

Construction 3.6. Let $\mathcal{H} = (G_H, H, \text{Inv})$ be a chameleon hash function, and $\mathcal{S}_c = (G_c, \text{Sign}_c, \text{Ver}_c)$ a signature scheme. Let \mathcal{Q} and \mathcal{R} be families of pairwise independent functions mapping messages to randomness used by Inv and Sign_c , respectively. We define a new signature scheme $\mathcal{S} = (G, \text{Sign}, \text{Ver})$ where:

$$G(\lambda) : (\text{sk}_H, \text{pk}_H) \stackrel{R}{\leftarrow} G_H(\lambda), (\text{sk}_c, \text{pk}_c) \stackrel{R}{\leftarrow} G_c(\lambda) \\ \text{output sk} = (\text{pk}_H, \text{sk}_c), \text{pk} = (\text{pk}_H, \text{pk}_c)$$

$$\text{Sign}((\text{pk}_H, \text{sk}_c), m) : Q \stackrel{R}{\leftarrow} \mathcal{Q}, R \stackrel{R}{\leftarrow} \mathcal{R} \\ r \leftarrow \text{Sample}(\lambda; R(m)), s \leftarrow Q(m), h \leftarrow H(\text{pk}_H, m, r) \\ \sigma \leftarrow \text{Sign}(\text{pk}_c, h; s), \text{output } (r, \sigma)$$

$$\text{Ver}((\text{pk}_H, \text{pk}_c), m, (r, \sigma)) : h \leftarrow H(\text{pk}_H, m, r), \text{output Ver}(\text{pk}_c, h, \sigma)$$

We note that the chameleon secret key is not used in Construction 3.6, though it will be used in the security proof. Classically, this method of hashing with a chameleon hash and then signing converts any non-adaptively secure scheme into an adaptive one. We show that the resulting scheme is actually secure against an adaptive *quantum* chosen message attack.

Theorem 3.7. *If \mathcal{S}_c is weakly (resp. strongly) EUF-RMA secure and \mathcal{H} is a secure chameleon hash function, then \mathcal{S} in Construction 3.6 is weakly (resp. strongly) EUF-qCMA secure. Moreover, if \mathcal{S}_c is only one-time secure, then \mathcal{S} is also one-time secure.*

Theorem 3.7 shows that we can take a classically EUF-RMA secure signature scheme, combine it with a chameleon hash, and obtain a quantum-secure signature scheme. In particular, the following constructions will be quantum secure, assuming SIS is hard:

- The signature scheme of Cash et al. [CHKP10], which combines their chameleon hash function with an EUF-RMA secure signature scheme.
- A modification of the signature scheme of Agrawal, Boneh, and Boyen [ABB10], where we hash the message using a chameleon hash before applying the signature.

We now prove Theorem 3.7:

Proof. We first sketch the proof idea: given an \mathcal{S}_c signature σ on a random hash h , we can construct an \mathcal{S} signature on any given message m by using the chameleon secret key sk_H to compute a randomness r such that $\text{H}(\text{pk}_H, m, r) = h$, and outputting the signature (r, σ) . This fact yields a simple proof of classical chosen message security. When quantum queries are considered, however, we need to sign an exponential sized set of messages, so using this technique directly would require signing an exponential number of random hashes. We can use small-range distributions and Lemma 2.4 to reduce the number of signed hashes to a polynomial. The problem is that the number of hashes signed is still a very large polynomial, whereas the number of signatures produced by our adversary is only $q + 1$, so we cannot rely on the pigeon-hole principle to argue that one of the \mathcal{S} forgeries is in fact a \mathcal{S}_c forgery. We instead pick a random query, and perform a particular partial measurement on that query superposition, to obtain one of polynomially many outcomes. Lemma 2.1 shows that the adversary's advantage is reduced by only a polynomial factor. Our measurement is performed in such a way that the adversary only sees a single signature for that query. Since the adversary produced more forgeries than queries, two of the forgeries must correspond to the same query. If this is the query that we measured, then one of the \mathcal{S} forgeries is actually a forgery for \mathcal{S}_c .

We now give the complete proof. There are four variants to the theorem (one-time vs many time, strong vs weak). We will prove the many-time strong security variant, the other proofs being similar. Let A be an adversary breaking the EUF-qCMA security of \mathcal{S} in Construction 3.6 with non-negligible probability ϵ . We prove security through a sequence of games.

Game 0. This is the standard attack experiment, where A receives pk_c and pk_H , and is allowed to make a polynomial number of quantum chosen message queries. For query i , the challenger produces pairwise independent functions $R^{(i)}$ and $Q^{(i)}$, and responds to each message in the query superposition as follows:

- Let $r_m^{(i)} = \text{Sample}(\lambda; R^{(i)}(m))$ and $s_m^{(i)} = Q^{(i)}(m)$.

- Compute $h_m^{(i)} = \mathsf{H}(m, r_m^{(i)})$
- Compute $\sigma_m^{(i)} = \mathsf{Sign}_c(\mathsf{sk}_c, h_m^{(i)}; s_m^{(i)})$
- Respond with the signature $(r_m^{(i)}, \sigma_m^{(i)})$.

In the end, A must produce $q+1$ distinct triples $(m_k^*, r_k^*, \sigma_k^*)$ such that $\mathsf{Ver}(\mathsf{pk}_c, \mathsf{H}(\mathsf{pk}_H, m_k^*, r_k^*), \sigma_k^*)$ accepts. By definition, A wins with probability ϵ , which is non-negligible. Therefore, there is some polynomial $p(\lambda)$ such that $p(\lambda) > 1/\epsilon(\lambda)$ for infinitely-many λ .

Game 1. We make two modifications: first, we choose $R^{(i)}$ and $Q^{(i)}$ as truly random functions, which amounts to generating $r_m^{(i)} \leftarrow \mathsf{Sample}(\lambda)$ and picking $s_m^{(i)}$ at random for each i, m . The adversary cannot tell the difference according to Lemma 2.2. Second, we modify the conditions in which A wins by requiring that no two (m_k^*, r_k^*) pairs form a collision for H . The security of \mathcal{H} implies that A succeeds in Game 1 with probability at least $\epsilon - \mathsf{negl}$.

Game 2. Generate $s_m^{(i)}$ as before, but now draw $h_m^{(i)}$ uniformly at random, as we as randomness $t_m^{(i)}$. We will sample $r_m^{(i)}$ from the set of randomness making $\mathsf{H}(\mathsf{pk}, m, r_m^{(i)}) = h_m^{(i)}$. That is, let $r_m^{(i)} = \mathsf{Inv}(\mathsf{sk}, h_m^{(i)}, m; t_m^{(i)})$. The only difference from A 's perspective is the distribution of the $r_m^{(i)}$ values. For each m , the distribution of $r_m^{(i)}$ is negligibly-close to that of Game 1, so the oracles mapping m to $r_m^{(i)}$ are indistinguishable from those in Game 1 by Lemma 2.5. Therefore, the success probability is at least $\epsilon - \mathsf{negl}$.

Game 3. Let $\ell = 2C_0qp$ where C_0 is the constant from Lemma 2.4. At the beginning of the game, for $i = 1, \dots, q$ and $j = 1, \dots, \ell$, sample values $\hat{h}_j^{(i)}$ and let $\hat{\sigma}_j^{(i)} = \mathsf{Sign}_c(\mathsf{sk}_c, \hat{h}_j^{(i)})$. Also pick q random functions O_i mapping m to $[\ell]$. Then let $h_m^{(i)} = \hat{h}_{O_i(m)}^{(i)}$ and $\sigma_m^{(i)} = \hat{\sigma}_{O_i(m)}^{(i)}$. Let T_i be random functions, and let $t_m^{(i)} = T_i(m)$. The only difference between Game 2 and Game 3 is that the $h_m^{(i)}$ and $\sigma_m^{(i)}$ values were generated by q small-range distributions on ℓ samples. Each of the small-range distributions is only queried once, so Lemma 2.4 implies that the success probability is still at least $\epsilon - \mathsf{negl} - 1/2p$.

Game 4. Let the O_i and T_i be pairwise independent functions. The adversary cannot tell the difference.

Notice that Game 4 can now be simulated efficiently, and A wins in this game with probability $\epsilon - \mathsf{negl} - 1/2p$. Let $h_k^* = \mathsf{H}(\mathsf{pk}, m_k^*, r_k^*)$ be the hashes of the forgeries. Since we have no collisions in H , the pairs (h_k^*, σ_k^*) are distinct. Let $\mathcal{H}^{(i)} = \{\hat{h}_j^{(i)}\}$ and \mathcal{H} be the union of the $\mathcal{H}^{(i)}$. There are two possibilities:

- At least one of the h_k^* is not in \mathcal{H} , or two of them are equal. In this case, we can obtain a forger B_0 for \mathcal{S}_c , which is given pk_c and simulates Game 4 exactly: To generate the $(\hat{h}_j^{(i)}, \hat{\sigma}_j^{(i)})$ pairs, B_0 asks its own challenger for signatures on $q\ell$ random messages. When A responds with forgeries $(m_k^*, r_k^*, \sigma_k^*)$, B_0 computes $h_k^* = \mathsf{H}(\mathsf{pk}_H, m_k^*, r_k^*)$, and finds the k value such that $h_k^* \notin \mathcal{H}$, or the k_0, k_1 such that $h_{k_0}^* = h_{k_1}^*$. In the latter case, one of the $\sigma_{k_b}^*$ was not the result of a signing query, so let $k = k_b$. It then outputs the pair (h_k^*, σ_k^*) . Then B_0 never received

the signature σ_k^* on h_k^* , so this is a valid forgery. Therefore, this event happens with negligible probability.

- All of the h_k^* values distinct and lie in \mathcal{H} . In this case, there is some i such that two h_k^* values are in $\mathcal{H}^{(i)}$ for the same i . Notice that this event happens, and all the forgeries are valid, with probability $\epsilon - \text{negl} - 1/2p$.

Game 5. Now we pick a random query i^* add a check that all the h_k^* values lie in \mathcal{H} , and that two of them are distinct and lie in $\mathcal{H}^{(i^*)}$. Without loss of generality, assume these two h^* values are h_0^* and h_1^* . A then wins in this game with probability $\epsilon/q - \text{negl} - 1/2pq$. Let j_b^* be the j such that $h_b = \hat{h}_{j_b^*}^{(i^*)}$ for $b = 0, 1$.

Game 6. On query i^* , measure the value of $O_i(m)$, to get a value j^* . O_i takes values in $[\ell]$, so the adversary's success probability is still at least $\epsilon/q\ell - \text{negl} - 1/2pq\ell$. Notice now that for query i^* , the challenger only needs to sign $\hat{h}_{j^*}^{(i^*)}$, and therefore, one of the $h_b = \hat{h}_{j_b^*}^{(i^*)}$ values was never signed.

Game 7. Now guess at the beginning of the game the value of j^* . For $j \neq j^*$, generate $\hat{h}_j^{(i^*)}$ randomly and do not sign them. For the rest of the $\hat{h}_j^{(i)}$, query the RMA challenger for \mathcal{S}_c for random message/signature pairs. Simulate game 6, except abort if we did not guess j^* correctly. The adversary still wins with probability at least $\epsilon/q\ell^2 - \text{negl} - 1/4pq\ell^2$, and if it does, we now have a forgery on h_b^* for some b . The security of \mathcal{S}_c implies therefore that $\epsilon/q\ell^2 - \text{negl} - 1/2pq\ell^2$ is negligible. Thus $\epsilon - 1/2p$ is negligible. Since $\epsilon > 1/p$ infinitely often, we then have $1/2p < \text{negl}$ infinitely often, a contradiction. Therefore, ϵ is negligible. \square

We note that for one-time security, this security reduction signs only a single message, so we only need to rely on the one-time security of \mathcal{S}_c .

3.2 Signatures in the Quantum Random Oracle Model

In this section, we demonstrate a generic conversion from a classical signature scheme to one that is secure against an adaptive *quantum* chosen message attack in the quantum random oracle model. In the quantum random oracle model, all parties can issue quantum queries to a random function H . The idea of our construction is simple: use the random oracle to hash the message along with a random salt, and send the signature on the hash, together with the salt:

Construction 3.8. Let $\mathcal{S}_c = (\mathbf{G}_c, \text{Sign}_c, \text{Ver}_c)$ be a signature scheme, H be a hash function, and \mathcal{Q} be a family of pairwise independent functions mapping messages to the randomness used by Sign_c , and k some polynomial in λ . Define $\mathcal{S} = (\mathbf{G}, \text{Sign}, \text{Ver})$ where:

$$\mathbf{G}(\lambda) = \mathbf{G}_c(\lambda)$$

$$\begin{aligned} \text{Sign}(\text{sk}, m) : Q \xleftarrow{R} \mathcal{Q}, r \xleftarrow{R} \{0, 1\}^k \\ s \leftarrow Q(m), h \leftarrow H(m, r), \sigma \leftarrow \text{Sign}_c(\text{sk}, h; s), \text{ output } (r, \sigma) \end{aligned}$$

$$\text{Ver}(\text{pk}, m, (r, \sigma)) : h \leftarrow H(m, r), \text{ output } \text{Ver}_c(\text{pk}, h, \sigma)$$

We note that Construction 3.8 is similar to Construction 3.6: instead of the chameleon hash $H(\text{pk}, \cdot, \cdot)$ we have a random oracle $H(\cdot, \cdot)$, and instead of generating a different r for each message in the superposition, we just generate a single r for the entire superposition. We can achieve security for Construction 3.8, assuming only a very weak form of security for \mathcal{S}_c , namely, universal unforgeability under a random message attack (UUF-RMA security):

Theorem 3.9. *If \mathcal{S}_c is strongly (resp. weakly) UUF-RMA secure, then \mathcal{S} in Construction 3.8 is strongly (resp. weakly) EUF-qCMA secure in the quantum random oracle model. Moreover, if \mathcal{S}_c is only one-time secure, then \mathcal{S} is also one-time secure.*

Before proving Theorem 3.9, we explain how to realize the strong UUF-RMA notion of security. We note that any strongly EUF-RMA or EUF-CMA secure signature scheme satisfies this security notion. We also note that some weaker primitives do as well. The first is pre-image sampleable functions, defined by Gentry et al. [GPV08]:

Definition 3.10 (PSF). *A pre-image sampleable trapdoor function (PSF) is a tuple of algorithms $\text{PSF} = (\text{G}, \text{Sample}, F, F^{-1})$ with the following properties:*

- $\text{G}(\lambda)$ generates a secret/public key pair (sk, pk) .
- $F(\text{pk}, \cdot)$ is a function from set \mathcal{X}_λ to set \mathcal{Y}_λ .
- $\text{Sample}(\lambda)$ samples an x from some on \mathcal{X}_λ , such that $F(\text{pk}, x)$ is uniform over \mathcal{Y}_λ .
- $F^{-1}(\text{sk}, y)$ takes an image $y \in \mathcal{Y}_\lambda$, and outputs an x such that $F(\text{pk}, x) = y$, and x is distributed negligibly-close to $\text{Sample}(\lambda)$ conditioned on $F(\text{pk}, x) = y$.

The two general notions of security we are interested in for PSFs are one-wayness and collision resistance. If we let $\text{Sign}(\text{sk}, m) = F^{-1}(\text{sk}, m)$ and $\text{Ver}(\text{pk}, m, \sigma) = F(\text{pk}, \sigma) == m$, then one-wayness plus collision resistance implies strong UUF-RMA security.

Corollary 3.11. *If PSF is a collision resistant and one-way PSF, then Construction 3.8 instantiated with PSF is strongly EUF-qCMA secure in the quantum random oracle model.*

Gentry et al. [GPV08] show how to construct a PSF that is collision-resistant and one-way under the assumption that SIS is hard. Therefore, we can construct efficient signatures in the quantum random oracle model based on SIS. At the end of this section, we also show that the basic GPV signature scheme is secure in the quantum random oracle model, though the proof is very different.

A trapdoor permutation is a PSF where D_λ is the uniform distribution and $F(\text{pk}, \cdot)$ is bijective. Trapdoor permutations are trivially collision resistant, since they have no collisions.

Corollary 3.12. *If \mathcal{F} is a one-way trapdoor permutation, then Construction 3.8 instantiated with \mathcal{F} is strongly EUF-qCMA secure in the quantum random oracle model.*

We now sketch the proof of Theorem 3.9. The complete proof is in Appendix A.1.

Proof sketch. Given the similarities between Constructions 3.6 and 3.8, the proof is similar to that of Theorem 3.7. For classical security in the classical random oracle model, the adversary only sees a polynomial number of outputs of H . We can set these outputs to be exactly the messages produced by the \mathcal{S}_c challenger. Moreover, we can set the outputs in a way so that we can answer signing queries using the signatures provided by the \mathcal{S}_c challenger with non-negligible probability. For

quantum security in the quantum random oracle model, using this approach directly would require the \mathcal{S}_c challenger to output exponentially many random messages, and sign an exponential number of them. Similar to the proof of Theorem 3.7, we can overcome this difficulty using small-range distributions. However, now the number of signatures received from the \mathcal{S}_c challenger is a large polynomial, whereas the adversary only produces $q + 1$ \mathcal{S} forgeries. To show that one of the forgeries still corresponds to an \mathcal{S}_c forgery, we perform a partial measurement on one of the queries, so that the adversary only sees a single signature for that query. Since the adversary produced $q + 1$ forgeries, two of them must correspond to the same query, so one of the \mathcal{S} forgeries must actually be an \mathcal{S}_c forgery. \square

Before completing this section, we show that the basic deterministic GPV signature scheme is secure. For completeness, we present the GPV signature scheme built from pre-image sampleable functions and PRFs, and prove its security:

Construction 3.13. Let $\text{PSF} = (\mathbf{G}_{psf}, \text{Sample}, F, F^{-1})$ be a pre-image sampleable function, PRF be a pseudorandom function, and H a hash function. Let $\mathcal{S} = (\mathbf{G}, \text{Sign}, \text{Ver})$ where

$$\mathbf{G}(\lambda) : (\text{sk}', \text{pk}') \xleftarrow{R} \mathbf{G}_{psf}(\lambda), k \xleftarrow{R} \{0, 1\}^\lambda \\ \text{output sk} = (\text{sk}', k), \text{pk} = \text{pk}'$$

$$\text{Sign}((\text{sk}, k), m) : r \leftarrow \text{PRF}(k, m) \ h \leftarrow H(m), \text{ output } \sigma = F^{-1}(\text{sk}, h; r)$$

$$\text{Ver}(\text{pk}, m, \sigma) : h \leftarrow H(m), h' \leftarrow F(\text{pk}, \sigma), \text{ accept if and only if } h = h'$$

We say that PSF has large pre-image min-entropy if, for all pk ,

$$\max_{y \in \mathcal{Y}} \Pr[x \leftarrow \text{Sample}(\lambda) : F(\text{pk}, x) = y] < 2^{-\omega(\log \lambda)}$$

We note that the PSF given by Gentry et al. [GPV08] has large pre-image min-entropy.

Theorem 3.14. *If PSF is collision resistant and has large pre-image min-entropy, then \mathcal{S} from Construction 3.13 is EUF-qCMA secure.*

Proof. We prove security via a sequence of games:

Game 0. This is the standard security game. The adversary wins with probability ϵ .

Game 1. Replace PRF with a truly random function. The security of PRF implies that the adversary wins with probability at least $\epsilon - \text{negl}$.

Game2. We change the way we answer signing queries and oracle queries as follows: Pick a random function J that maps messages to the randomness used by $\text{Sample}(\lambda)$. We implement the signing oracle as $S(m) = \text{Sample}(\lambda; J(m))$. That is, signatures are random samples from D_λ , where the randomness used in the sampling is obtained by $J(m)$. We implement the random oracle as $H(m) = F(\text{pk}, S(m))$. The adversary wins if he can produce $q + 1$ (m_i, σ_i) pairs where $H(m_i) = F(\text{pk}, \sigma_i)$. This corresponds to $F(\text{pk}, S(m_i)) = F(\text{pk}, \sigma_i)$. In other words, $S(m_i)$ and σ_i form a collision. By the collision resistance of PSF, we must have $S(m_i) = \sigma_i$ for all i , except with negligible probability. This means that we make q queries to the oracle S and a polynomial number

of queries to the oracle $F(\text{pk}, S(\cdot))$, and output $q + 1$ input/output pairs of S with probability $\epsilon - \text{negl}$.

Game 2 satisfies the conditions of Lemma 2.6 with $H = S$, $F(\cdot) = F(\text{pk}, \cdot)$, and $H_\infty = \omega(\log \lambda)$. Therefore, the probability A wins in Game 2 is at most $(q + 1) / \lfloor 2^{H_\infty} \rfloor < (q + 1) 2^{-\omega(\log \lambda)}$, which is negligible. Therefore, A wins in Game 0 with negligible probability, as desired. \square

3.3 Signatures from Generic Assumptions

In this section, we show how to construct signatures from generic assumptions. We first construct one-time signatures from one-way functions using the basic Lamport construction. We then expand the message space to handle arbitrary-length messages using collision resistance, and finally plug these one-time signatures into the Merkle signature scheme. The end result is a signature scheme whose quantum security relies only on the existence of collision-resistant functions:

Theorem 3.15. *If there exists a collision-resistant hash function, then there exists a strongly EUF-qCMA secure signature scheme.*

We now give the basic Lamport scheme and prove its security:

Construction 3.16. *Let F be a one-way function. We define the following signature scheme for n -bit messages:*

$$\begin{aligned} \text{G}(\lambda) : & \text{for each } i \in [n], b \in \{0, 1\}, x_{i,b} \xleftarrow{R} \{0, 1\}^\lambda, y_{i,b} \leftarrow F(x_i) \\ & \text{output } \text{sk} = (x_{i,b})_{i \in [n], b \in \{0, 1\}}, \text{pk} = (y_{i,b})_{i \in [n], b \in \{0, 1\}} \end{aligned}$$

$$\begin{aligned} \text{Sign}(\text{sk}, m) : & \text{write } \text{sk} = (x_{i,b})_{i \in [n], b \in \{0, 1\}} \\ & \text{output } (x_{i, m_i})_{i \in [n]} \end{aligned}$$

$$\begin{aligned} \text{Ver}(\text{pk}, m, \sigma) : & \text{write } \text{pk} = (y_{i,b})_{i \in [n], b \in \{0, 1\}}, \sigma = (x'_i)_{i \in [n]} \\ & \text{accept if and only if } F(x'_i) = y_{i, m_i} \text{ for all } i \in [n] \end{aligned}$$

Theorem 3.17. *If F is one-way (resp. second pre-image resistant), then the Lamport signature scheme built from F is weakly (resp. strongly) one-time EUF-CMA secure.*

Proof. We prove the weak security case; the strong security case is almost identical. Let A be an adversary that makes a single quantum query to Sign and outputs a pair of valid message/signature pairs for different messages with probability ϵ . We prove security through a sequence of games.

Game 0. This is the standard attack game, where A wins with probability ϵ .

Game 1. Pick a random value $i^* \in [n]$. Abort if both messages in A 's forgery are the same for index i^* . A still wins with probability ϵ/n .

Game 2. For the quantum chosen message query, measure the bit i^* of the message superposition. Lemma 2.1 shows that A still wins with probability $\epsilon/2n$.

Game 3. At the beginning of the game, guess a bit b^* at random, and abort if the outcome of the measurement in Game 2 is b^* . A still wins with probability $\epsilon/4n$.

We can now describe an adversary B that inverts F . On input y , B guesses $i^* \in [n]$ and $b^* \in \{0, 1\}$, and sets $y_{i^*, b^*} = y$. For $(i, b) \neq (i^*, b^*)$, B picks $x_{i, b}$ at random, and lets $y_{i, b} = F(x_{i, b})$. Now B simulates Game 3. With probability at least $\epsilon/4n$, A produces valid forgeries whose messages differ on bit i^* . This means A produces pre-images $x'_{i^*, 0}, x'_{i^*, 1}$ for $y_{i^*, 0}, y_{i^*, 1}$. B outputs x'_{i^*, b^*} , which is a valid pre-image for $y_{i^*, b^*} = y$. □

The signatures from Construction 3.16 have public keys that are much longer than the messages being signed. In order to use Lamport signatures in the Merkle signature scheme, we need to be able to sign much larger messages. In the classical setting, it is possible to expand the message space using target collision resistant functions. These can in turn be built from one-way functions, showing that classical signatures can be built from the minimal assumption of one-way functions. Unfortunately, the notion of target collision resistance no longer makes sense in the quantum setting, and we therefore have to resort to collision resistance. Therefore, we can build one-time signatures for arbitrary-length messages assuming only collision resistance.

Now we show how to use such signatures to build many-time signatures. We will have a tree of depth d , where each non-leaf node contains a pair of private/public key pairs for the one-time signature scheme, one for each child. The private/public keys for the system will be the keys for the root. To sign a message, a random leaf node is chosen. For each non-leaf node in the path from root to leaf, sign the node's public keys with the corresponding secret key of the parent. Then use the correct secret key from the leaf's parent to sign the message. This tree is exponential in size, so we will use a PRF to generate the keys. In more detail:

Construction 3.18. Let $\mathcal{S}_{ot} = (\mathsf{G}_{ot}, \mathsf{Sign}_{ot}, \mathsf{Ver}_{ot})$ be a one-time signature scheme. Also let F be a secure PRF. The stateless Merkle signature scheme $\mathcal{S} = (\mathsf{G}, \mathsf{SignVer})$ is defined as follows:

- $\mathsf{G}(\lambda)$: run G_{ot} twice to get two secret/public key pairs $(\mathsf{sk}_b, \mathsf{pk}_b)$ for $b = 0, 1$. Also choose a random λ -bit string k . The secret key is $\mathsf{sk} = (\mathsf{sk}_0, \mathsf{sk}_1, k)$ and the public key is $\mathsf{pk} = (\mathsf{pk}_0, \mathsf{pk}_1)$.
- $\mathsf{Sign}(\mathsf{sk}, m)$: to sign a message m , first pick a random bit string \mathbf{a} in $\{0, 1\}^d$. Then for $i = 1, \dots, d - 1$,

- Let $(r_{(\mathbf{a}_{[1, i]}, 0)}, r_{(\mathbf{a}_{[1, i]}, 1)}, s_{\mathbf{a}_{[1, i]}}) = F(k, \mathbf{a}_{[1, i]})$.
- Let $(\mathsf{sk}_{(\mathbf{a}_{[1, i]}, b)}, \mathsf{pk}_{(\mathbf{a}_{[1, i]}, b)}) = \mathsf{G}_{ot}(\lambda; r_{(\mathbf{a}_{[1, i]}, b)})$ for $b = 0, 1$.
- Let $\sigma_{\mathbf{a}_{[1, i]}} = \mathsf{Sign}_{ot}(\mathsf{sk}_{\mathbf{a}_{[1, i]}}, (\mathsf{pk}_{(\mathbf{a}_{[1, i]}, 0)}, \mathsf{pk}_{(\mathbf{a}_{[1, i]}, 1)}); s_{\mathbf{a}_{[1, i]}})$

Let $\Sigma_{\mathbf{a}} = (\mathbf{a}, (\mathsf{pk}_{(\mathbf{a}_1, 0)}, \mathsf{pk}_{(\mathbf{a}_1, 1)}, \sigma_{\mathbf{a}_1}), \dots, (\mathsf{pk}_{(\mathbf{a}_{[1, d-1]}, 0)}, \mathsf{pk}_{(\mathbf{a}_{[1, d-1]}, 1)}))$, and let $\sigma_{\mathbf{a}}(m) = \mathsf{Sign}_{ot}(\mathsf{sk}_{\mathbf{a}}, m)$. Output the signature $(\Sigma_{\mathbf{a}}, \sigma_{\mathbf{a}}(m))$.

- $\mathsf{Ver}(\mathsf{pk}, m, \Sigma, \sigma)$: parse Σ as $(\mathbf{a}, (\mathsf{pk}_{(\mathbf{a}_1, 0)}, \mathsf{pk}_{(\mathbf{a}_1, 1)}, \sigma_{\mathbf{a}_1}), \dots, (\mathsf{pk}_{(\mathbf{a}_{[1, d-1]}, 0)}, \mathsf{pk}_{(\mathbf{a}_{[1, d-1]}, 1)}))$. For $i = 1, \dots, d - 1$,
- If $\mathsf{Ver}_{ot}(\mathsf{pk}_{\mathbf{a}_{[1, i]}}, (\mathsf{pk}_{(\mathbf{a}_{[1, i]}, 0)}, \mathsf{pk}_{(\mathbf{a}_{[1, i]}, 1)}))$ rejects, then reject and stop.

Then output the output of $\mathsf{Ver}_{ot}(\mathsf{pk}_{\mathbf{a}}, m, \sigma)$.

We note that if we allow state, we can pick the random bit string \mathbf{a} incrementally for each query. Then we can actually save the $\text{sk}_{(\mathbf{a}_{[1,i]}, b)}$, $\text{pk}_{(\mathbf{a}_{[1,i]}, b)}$, $\sigma_{\mathbf{a}_{[1,i]}}$ values until we do not need them any more, and remove the need for a PRF to generate randomness. In this way, we obtain the stateful Merkle signature scheme.

Theorem 3.19. *If \mathcal{S}_{ot} is weakly (resp. strongly) one-time EUF-qCMA secure, then both the stateless and stateful Merkle Signatures built from \mathcal{S}_{ot} are weakly (resp. strongly) EUF-qCMA secure.*

Proof sketch. The proof is very similar to the classical proof. Notice that each secret key in all but the bottom level are used to sign exactly one message: the pair of public keys in the corresponding child. Moreover, the secret keys on the bottom level are used to answer only one (potentially quantum) signature query. Therefore, the security of the one-time signature scheme implies that no adversary can forge messages for the Merkle signature scheme. For completeness, we give the complete proof in Appendix A.2. \square

4 Encryption Schemes

We now turn to encryption schemes where we first discuss an adequate notion of security under quantum queries. In what follows, we will discuss symmetric key schemes; the discussion for public key schemes is similar. We start by developing a notion of CPA security where encryption queries are allowed to be quantum. Since defining an attainable definition is non-trivial we present a few alternatives and converge on a reasonable definition. Once we arrive at a suitable definition for CPA security, we will generalize it to a corresponding CCA security definition.

Definition 4.1. *A symmetric key encryption scheme $\mathcal{E} = (\text{Enc}, \text{Dec})$ is indistinguishable under a fully quantum chosen plaintext attack (IND-fqCPA secure) if no efficient adversary A can win the following game, except with probability at most $1/2 + \text{negl}$:*

Key Gen *The challenger picks a random key k and a random bit b .*

Encryption Queries *A is allowed to make chosen message queries on superpositions of message pairs. For each such query, the challenger chooses randomness r , and encrypts the appropriate message in each pair using r as randomness:*

$$\sum_{m_0, m_1, c} \psi_{m_0, m_1, c} |m_0, m_1, c\rangle \quad \longrightarrow \quad \sum_{m_0, m_1, c} \psi_{m_0, m_1, c} |m_0, m_1, c \oplus \text{Enc}(k, m_b; r)\rangle$$

Guess *A produces a bit b' , and wins if $b = b'$.*

Definition 4.1 captures a scheme where we can encrypt a superposition of messages by encrypting each message in the superposition separately, and no efficient adversary can learn anything about the plaintext superposition. Unfortunately, this definition is not achievable:

Theorem 4.2. *No encryption scheme \mathcal{E} satisfies the security notion of Definition 4.1.*

Proof. We construct a generic adversary A . A prepares three registers: two plaintext registers and a ciphertext register. A puts a uniform superposition of all messages in the first register, and 0 in the second plaintext and ciphertext registers. A submits these three registers as a chosen message query. If $b = 0$, the ciphertext register will contain the encryptions of the messages in the

superposition. If $b = 1$, it will contain the encryption of 0. A then measures the ciphertext register. If $b = 0$, the resulting state will be the purely classical state $(m, 0, \text{Enc}(k, m))$ for a random message m . If $b = 1$, the measurement does nothing, so the first register still contains a superposition of all messages. A now performs the quantum Fourier transform to the first message register and measures. If $b = 0$, the transform will place a uniform superposition of all messages in the first register, and measuring will give a random message. If $b = 1$, the transform will place 0 in the first register. Thus, A distinguishes $b = 0$ from $b = 1$ with probability exponentially-close to 1. \square

The problem with Definition 4.1 is that the message query is entangled with the ciphertext response, and this entanglement depends on which register gets encrypted. Another reasonable idea is to encrypt *both* registers, but flip which register each ciphertext is written to depending on the value of b :

Definition 4.3. *A symmetric key encryption scheme $\mathcal{E} = (\text{Enc}, \text{Dec})$ is indistinguishable under a fully quantum chosen left-right plaintext attack (IND-lrCPA secure) if no efficient adversary A can win in the following game, except with probability at most $1/2 + \text{negl}$:*

Key Gen *The challenger picks a random key k and a random bit b .*

Encryption Queries *A is allowed to make chosen message queries. For each such query, the challenger chooses randomness r_0, r_1 , and responds with the encryptions of both messages in the pair, but in an order determined by b :*

$$\sum_{m_0, m_1, c_1, c_2} \psi_{m_0, m_1, c_1, c_2} |m_0, m_1, c_1, c_2\rangle \quad \longrightarrow \quad \sum_{m_0, m_1, c_1, c_2} \psi_{m_0, m_1, c_1, c_2} |m_0, m_1, c_1 \oplus \text{Enc}(k, m_b; r_0), c_2 \oplus \text{Enc}(k, m_{1-b}; r_1)\rangle$$

Guess *A produces a bit b' , and wins if $b = b'$.*

Unfortunately, this definition turns out to be at least as strong as Definition 4.1, and so it is also unattainable:

Theorem 4.4. *No encryption scheme \mathcal{E} satisfies the notion of security in Definition 4.3. In particular, any encryption scheme that is secure in the sense of Definition 4.3 is also secure in the sense of Definition 4.1.*

Proof. Suppose we have an adversary A for Definition 4.1. We will convert it into an adversary B for Definition 4.3. B simulates A forwarding encryption queries as follows: When A makes an encryption query, B adds a second ciphertext register, and puts into it a uniform superposition over all ciphertexts. B then sends the resulting state to its challenger as its encryption query. The answer to this query does not affect the second ciphertext register, so B can uncompute it. B then passes the resulting state back to A . B perfectly simulates A 's view, and therefore B breaks the security of \mathcal{E} under Definition 4.3 \square

Our attempts to make the entire security game quantum lead to an adversary that can always win. Therefore, we must allow encryption queries to be classical. We do, however, wish to allow the adversary to encrypt superpositions of messages, but not have the response depend in any way on b . Therefore, we propose separating encryption queries into classical challenge queries and quantum encryption queries. This gives the following definition:

Definition 4.5. A symmetric key encryption scheme $\mathcal{E} = (\text{Enc}, \text{Dec})$ is indistinguishable under a quantum chosen message attack (IND-qCPA secure) if no efficient adversary A can win in the following game, except with probability at most $1/2 + \text{negl}$:

Key Gen The challenger picks a random key k and a random bit b .

Queries A is allowed to make two types of queries:

Challenge queries A sends two messages m_0, m_1 , to which the challenger responds with $c^* = \text{Enc}(k, m_b)$.

Encryption queries For each such query, the challenger chooses randomness r , and encrypts each message in the superposition using r as randomness:

$$\sum_{m,c} \psi_{m,c} |m, c\rangle \quad \longrightarrow \quad \sum_{m,c} \psi_{m,c} |m, c \oplus \text{Enc}(k, m; r)\rangle$$

Guess A produces a bit b' , and wins if $b = b'$.

This definition has another advantage: since challenge queries are classical, when we move to CCA security, we can actually check if a ciphertext was the result of a challenge query. This gives us the following notion of CCA security:

Definition 4.6. A symmetric key encryption scheme $\mathcal{E} = (\text{Enc}, \text{Dec})$ is indistinguishable under a quantum chosen message attack (IND-qCCA secure) if no efficient adversary A can win in the following game, except with probability at most $1/2 + \text{negl}$:

Key Gen The challenger picks a random key k and a random bit b . It also creates a list \mathcal{C} which will store challenger ciphertexts.

Queries A is allowed to make three types of queries:

Challenge queries A sends two messages m_0, m_1 , to which the challenger responds with $c^* = \text{Enc}(k, m_b)$. The challenger also adds c^* to \mathcal{C} .

Encryption queries For each such query, the challenger chooses randomness r , and encrypts each message in the superposition using r as randomness:

$$\sum_{m,c} \psi_{m,c} |m, c\rangle \quad \longrightarrow \quad \sum_{m,c} \psi_{m,c} |m, c \oplus \text{Enc}(k, m; r)\rangle$$

Decryption queries For each such query, the challenger decrypts all ciphertexts in the superposition, except those that were the result of a challenge query:

$$\sum_{c,m} \psi_{c,m} |c, m\rangle \quad \longrightarrow \quad \sum_{c,m} \psi_{c,m} |c, m \oplus f(c)\rangle$$

where

$$f(c) = \begin{cases} \perp & \text{if } c \in \mathcal{C} \\ \text{Dec}(k, c) & \text{otherwise} \end{cases}$$

Guess A produces a bit b' , and wins if $b = b'$.

It is straightforward to modify the above definition for public key encryption schemes.

4.1 Symmetric CCA Security

In this section, we construct symmetric-key CCA secure encryption. We will follow the encrypt-then-MAC paradigm. Ideally, we would like to show that encrypt-then-MAC, when instantiated with any IND-qCPA-secure encryption scheme and any EUF-qCMA MAC, would be CCA secure. However, it is not obvious how to prove security, as the reduction algorithm has no way to tell which ciphertexts the adversary received as the result of an encryption query, and no way to decrypt the ciphertexts if it has received them. To remedy these problems, we choose a specific encryption scheme and MAC and leave the general security proof as an open question. The encryption scheme allows us to efficiently check if the adversary has seen a particular ciphertext as a result of an encryption query, and to decrypt in this case. The construction is as follows:

Construction 4.7. *Let F and G be pseudorandom functions. We construct the following encryption scheme $\mathcal{E} = (\text{Enc}, \text{Dec})$ where:*

$$\begin{aligned} \text{Enc}((k_1, k_2), m) : & r \xleftarrow{R} \{0, 1\}^\lambda \\ & c_1 \leftarrow F(k_1, r) \oplus m, \quad c_2 \leftarrow G(k_2, (r, m)) \\ & \text{output } (r, c_1, c_2) \\ \text{Dec}((k_1, k_2), (r, c_1, c_2)) : & m \leftarrow c_1 \oplus F(k_1, r), \quad c'_2 \leftarrow G(k_2, (r, m)) \\ & \text{if } c_2 \neq c'_2, \text{ output } \perp \\ & \text{otherwise, output } m \end{aligned}$$

For security, we require F to be a classically secure PRF, and G to be quantum secure — secure against queries on a superposition of inputs. Zhandry [Zha12b] shows how to construct PRFs meeting this strong notion of security.

Theorem 4.8. *If F and G are quantum-secure pseudorandom functions, then \mathcal{E} in Construction 4.7 is qCCA-secure.*

As demonstrated by Zhandry [Zha12b], quantum-secure pseudorandom functions can be built from any one-way function. Therefore, Theorem 4.8 shows that quantum chosen ciphertext security can be obtained from the minimal assumption that one-way functions exist. We now give the proof of Theorem 4.8:

Proof. We first sketch the proof: we can replace F and G with random functions and only negligibly affect the success probability. Since each encryption query receives a single r for the entire query superposition, we can answer any encryption query by making a single query to F on r . It is easy to check if a ciphertext (r', c_1, c_2) was computed during an encryption query: just check if $r = r'$. We can also decrypt such a ciphertext, since we have seen $F(k_1, r)$. Including $c_2 = G(k_2, (r, m))$ in the ciphertext guarantees with overwhelming probability that the adversary can only submit valid ciphertexts if they were ciphertexts received during an encryption query, so we might as well reject all ciphertexts (r', c_1, c_2) where r' was not the randomness used in any encryption query. Now, the value of m_b in the challenge query becomes perfectly hidden, which means that the distinguishing probability is 0.

We now give the complete security proof: assume we have an adversary A that breaks the indistinguishability of \mathcal{E} in Construction 4.7 with probability ϵ . We prove security through a sequence of games.

Game 0. This is the standard attack game where A makes q_e encryption queries which are answered using randomness values r_i , q_c challenge queries which are answered using randomness r_i^* , and q_d decryption queries. Let $(m_{i,0}^*, m_{i,1}^*)$ denote the i th challenger query, and (r_i^*, c_i^*, d_i^*) be the response.

Game 1. Replace F and G with truly random functions. That is, answer the i th encryption query by mapping m to $(r_i, F(r_i) \oplus m, G(r_i, m))$, the i th challenge query with $(r_i^*, F(r_i^*) \oplus m_{i,b}^*, G(r_i^*, m_{i,b}^*))$, and answer decryption queries accordingly. Since F and G are quantum-secure pseudorandom functions, the advantage of A in Game 1 is at least $\epsilon - \text{negl}$.

Game 2. Now we abort if there is a collision among any of the r_i or r_i^* . The probability of a collision is at most $(q_e + q_c)^2 / 2|\mathcal{R}|$ where \mathcal{R} is the randomness space. This quantity is negligible, so A 's advantage is still $\epsilon - \text{negl}$.

Notice that we can pick the r_i values and r_i^* value at the start of the game, and query F on these values. Let $\mathcal{T}_i = \{r_j : j \leq i\}$ and $\mathcal{T}_i^* = \{r_j^* : j \leq i\}$. Also let $\mathcal{T} = \mathcal{S}_{q_e}$ and $\mathcal{T}^* = \mathcal{S}_{q_c}^*$. Notice that at any point, A never gets to see $G(r, m)$ for any m if $r \notin \mathcal{T}_i \cup \mathcal{T}_j^*$ where i is the number of encryption queries made so far and j is the number of challenge queries made so far. Note also that A only gets to see $G(r_k^*, m)$ where $m = m_{k,b}^*$.

Game 3. For a decryption query on a superposition of ciphertexts (r, c, d) , let n_e be the number of encryption queries made so far and n_c the number of challenge queries. Check that $r \in \mathcal{T}_{n_e}$, and respond with \perp for that slot otherwise. We now consider the ciphertexts that would be accepted in Game 2 but rejected in Game 3. Such ciphertexts come in two forms:

- $r \in \mathcal{T}_{n_c}^*$: Then $r = r_i^*$ for some i . In order to not be rejected in Game 2, we must have $c \neq c_i^*$ or $d \neq d_i^*$. In the first case, (r, c, d) is an encryption of a message $m \neq m_{i,b}^*$, so the value of $G(r_i^*, m)$ is hidden to the adversary. Therefore, the probability (r, c, d) is a valid ciphertext is negligible. In the second case, (r, c, d) is an encryption of $m_{i,b}^*$, but then d is not a valid MAC, so decryption fails.
- $r \notin \mathcal{T}_{n_e} \cup \mathcal{T}_{n_c}^*$: Then the value of $G(r, m)$ is completely hidden from the adversary, so the probability d is a valid MAC is negligible.

Therefore, the probability of rejection for any ciphertext in Game 3 is only negligibly higher than that in Game 2. This means that with overwhelming probability, we only changed the decryption oracle on a negligible fraction of inputs, so A can only distinguish Games 2 and 3 with negligible probability. Therefore, A 's advantage is still $\epsilon - \text{negl}$.

Game 4. Now notice that F is never queried except on the points r_i and r_i^* . Therefore, at the start of the game, we can pick random values f_i and f_i^* to correspond to $F(r_i)$ and $F(r_i^*)$. We can also pick random values g_i^* that correspond to $G(r_i^*, m_{i,b}^*)$ (since we only query G on this point once). The adversary's view in this game is unchanged, so A 's advantage is at least $\epsilon - \text{negl}$.

Notice that we answer the i th challenge query with $(r_i^*, f_i^* \oplus m_{i,b}^*, g_i^*)$, and that the values of f_i^* and g_i^* are never used again. This means that $m_{i,b}^*$ is statistically hidden from the adversary. Therefore, A 's advantage in Game 4 is identically 0, so $\epsilon = \text{negl}$. \square

4.2 Public-key CCA Security

In this section, we construct CCA-secure signatures in the public-key setting. The basic idea is to first build a selectively secure identity-based encryption scheme — whose security can be based on the Learning With Errors (LWE) Problem — and then adapt the generic transformation to CCA-security to the quantum setting:

Let $\mathcal{E}_{ibe} = (\mathsf{G}_{ibe}, \mathsf{Enc}_{ibe}, \mathsf{Dec}_{ibe}, \mathsf{Ext})$ be an IBE scheme that is selectively secure against quantum queries. It is straightforward to show that the basic IBE scheme of Agrawal, Boneh, and Boyen [ABB10] meets this security notion assuming LWE is hard. Let $\mathcal{S} = (\mathsf{G}_s, \mathsf{Sign}, \mathsf{Ver})$ be a strongly EUF-CMA secure one-time signature scheme (quantum security is unnecessary). We now construct an encryption scheme using the generic transformation from IBE to CCA security due to Boneh et al. [BCHK04]:

Construction 4.9. $\mathcal{E} = (\mathsf{G}, \mathsf{Enc}, \mathsf{Dec})$ where

$$\begin{aligned} \mathsf{G}(\lambda) &: \mathsf{G}_{ibe}(\lambda) \\ \mathsf{Enc}(\mathsf{mpk}, m) &: (\mathsf{sk}, \mathsf{vk}) \leftarrow \mathsf{G}_s(\lambda) \\ & \quad c \leftarrow \mathsf{Enc}_{ibe}(\mathsf{mpk}, \mathsf{vk}, m), \sigma \leftarrow \mathsf{Sign}(\mathsf{sk}, c) \\ & \quad \text{output } (\mathsf{vk}, c, \sigma) \\ \mathsf{Dec}(\mathsf{msk}, (\mathsf{vk}, c, \sigma)) &: \text{if } \mathsf{Ver}(\mathsf{vk}, c, \sigma) \text{ rejects, output } \perp \\ & \quad \mathsf{sk}_{\mathsf{vk}} \leftarrow \mathsf{Ext}(\mathsf{msk}, \mathsf{vk}), m \leftarrow \mathsf{Dec}_{ibe}(\mathsf{sk}_{\mathsf{vk}}, c), \text{ output } m \end{aligned}$$

It is not difficult to adapt the classical security proof to the quantum setting, showing that the above construction achieves quantum CCA security:

Theorem 4.10. *If the LWE problem is hard for quantum computers, then there exists a public key encryption scheme that is IND-qCCA secure.*

References

- [ABB10] Shweta Agrawal, Dan Boneh, and Xavier Boyen. Efficient lattice (H)IBE in the standard model. *Proceedings of EUROCRYPT*, pages 1–40, 2010.
- [BCHK04] Dan Boneh, Ran Canetti, Shai Halevi, and Jonathan Katz. Chosen-Ciphertext Security from Identity-Based Encryption. *Proceedings of EUROCRYPT*, pages 1–31, 2004.
- [BDF⁺11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random Oracles in a Quantum World. In *Proceedings of ASIACRYPT*, 2011.
- [BHK⁺11] Gilles Brassard, Peter Høyer, Kassem Kalach, Marc Kaplan, Sophie Laplante, and Louis Salvail. Merkle Puzzles in a Quantum World. *Proceedings of CRYPTO*, pages 391–410, 2011.
- [BS08] Gilles Brassard and Louis Salvail. Quantum Merkle Puzzles. *Second International Conference on Quantum, Nano and Micro Technologies (ICQNM 2008)*, pages 76–79, February 2008.

- [BZ13] Dan Boneh and Mark Zhandry. Quantum-secure message authentication codes. In *Proceedings of Eurocrypt (to appear)*, 2013. Full version available at the Electronic Colloquium on Computational Complexity: <http://eccc.hpi-web.de/report/2012/136>.
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *Proceedings of FOCS*. IEEE, 2001.
- [CHKP10] David Cash, Dennis Hofheinz, Eike Kiltz, and Chris Peikert. Bonsai Trees, or How to Delegate a Lattice Basis. *Proceedings of EUROCRYPT*, pages 523–552, 2010.
- [DFNS11] Ivan Damgård, Jakob Funder, Jesper Buus Nielsen, and Louis Salvail. Superposition attacks on cryptographic protocols. *CoRR*, abs/1108.6313, 2011.
- [GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for Hard Lattices and New Cryptographic Constructions. *Proceedings of the 40th Annual ACM symposium on Theory of computing (STOC)*, page 197, 2008.
- [HSS11] Sean Hallgren, Adam Smith, and Fang Song. Classical cryptographic protocols in a quantum world. In *Proceedings of CRYPTO*, LNCS. Springer, 2011.
- [IBM12] IBM Research. IBM research advances device performance for quantum computing, Feb. 2012. <http://www-03.ibm.com/press/us/en/pressrelease/36901.wss>.
- [KR00] Hugo Krawczyk and Tal Rabin. Chameleon hashing and signatures. In *Proc. of NDSS*, pages 1–22, 2000.
- [Unr10] Dominique Unruh. Universally Composable Quantum Multi-Party Computation. *Proceedings of EUROCRYPT*, pages 486–505, 2010.
- [Zha12a] Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. In *Proceedings of CRYPTO*, 2012. Full version available at the Cryptology ePrint Archives: <http://eprint.iacr.org/2012/076/>.
- [Zha12b] Mark Zhandry. how to construct quantum random functions. In *Proceedings of FOCS*, 2012. Full version available at the Cryptology ePrint Archives: <http://eprint.iacr.org/2012/182/>.

A Signature Proofs

A.1 Proof of Theorem 3.9

Proof. Suppose we have an adversary A that breaks the security of \mathcal{S} from Construction 3.8. Let q_S be the number of signing queries made by A , and q_H be the number of hash queries (including those used in signing). We will prove security through a sequence of games.

Game 0. This is the standard attack game. A makes q quantum chosen message queries, and succeeds if it produces $q + 1$ valid message/signature tuples $(m_j^*, r_j^*, \sigma_j^*)$. Let r_i be the random value produced in the i th query, and Q_i be the pair-wise independent functions. A 's success probability is, by assumption, some non-negligible quantity ϵ . Then there is some polynomial $p(\lambda)$ such that $p(\lambda) > 1/\epsilon(\lambda)$ infinitely often.

Game 1. Replace the R_i with a truly random function, and abort if any of the r_i values are identical. Then the success probability is at least $\epsilon - q^2/2^{k+1} \geq \epsilon - \text{negl}$. Notice that if all the r_i are distinct, we can replace $Q_i(m)$ with $Q(m, r_i)$ for a random oracle Q that is fixed across all queries. That is, we sign the i th query with the oracle that maps m to $(r_i, \text{Sign}_c(\text{sk}, H(m, r_i); Q(m, r_i)))$. Notice that the function $H'(m, r) = (H(m, r), Q(m, r))$ is a random function.

Game 2. Let $\ell = 6C_0pq_H^3$. We now change how H' is generated. Pick three random oracles U, V and W , where the codomain of U and V is $[\ell]$, and let $H'(m, r) = W(U(m, V(r)), V(r))$. What this distribution represents is, for each $V(r)$ value, picking a random small-range function on ℓ samples. In essence, we have a small-range distribution on small-range distributions. A simple generalization of Lemma 2.4 shows that this is indistinguishable from Game 1 except with probability $C_0q_H^3/\ell = 1/3p$.

Game 3. Pick the r_i values up front, and let \mathcal{R} be the set of r_i values. Abort if $V(r_i) = V(r_j)$ for any $i \neq j$. We can assume without loss of generality that $V(r_i) = i$. The probability of this abort is at most $q_S^2/2\ell \leq 1/12C_0q_Hp < 1/3p$. Therefore, A wins in Game 3 with probability at least $\epsilon - 2/3p$.

The following modifications are indistinguishable to the adversary: before the start of the game, draw ℓ^2 different $\hat{h}_{i,j}$ values. Sign each of them with $i \leq q_S$ using \mathcal{S} to get $\hat{\sigma}_{i,j}$. Then let $H(m, r) = \hat{h}_{V(r), U(m, V(r))}$ and sign the i th query my mapping m to $(r_i, \hat{\sigma}_{i, U(m, i)})$. We can also generate V and U from $2q_H$ -wise independent functions, and the adversary cannot tell.

Game 4. Pick a random r_{i_0} from \mathcal{R} . Add the condition that if the r_j^* all lie in \mathcal{R} , that the two that are equal must be r_{i_0} . This condition is independent of the view of the adversary, so the adversary wins with probability at least $(\epsilon - 2/3p)/q_S$.

Game 5. Measure the value of $U(m, i_0)$ for the i_0 th query. The adversary still wins with probability at least $(\epsilon - 2/3p)/q_S\ell$.

Game 6. Pick a random $j_0 \in [\ell]$, and abort if the result of the measurement in Game 5 does not yield j_0 . We guess right with probability $1/\ell$, so the adversary still wins with probability at least $(\epsilon - 2/3p)/q_S\ell^2$. Now, if we succeed, we never need the values of $\hat{\sigma}_{i_0, j}$ except for $\hat{\sigma}_{i_0, j_0}$, so we don't need to ever sign the others.

We can now describe an adversary B that attacks the UUF-RMA security of \mathcal{S}_c . B simulates the entire Game 6, except for generating the $\hat{h}_{i,j}$ and $\hat{\sigma}_{i,j}$. For these, B asks its \mathcal{S} challenger for $q = (q_S - 1)\ell + 1$ random message/signature pairs, and $n = \ell^2 - q$ additional random messages. It assigns the q message/signature pairs to $\hat{h}_{i,j}$ and $\hat{\sigma}_{i,j}$ for $i \in [q_S] \setminus \{i_0\}$ and $\hat{\sigma}_{i_0, j_0}$. The rest of the $\hat{h}_{i,j}$ it sets to the n additional messages. When A outputs its $q_S + 1$ forgery candidates, there are several possibilities:

- $r_{j_1}^*$ lies outside \mathcal{R} for some j_1 . In this case, since there are no collisions among the (m_j^*, r_j^*) , $h_{j_1}^* = H(m_{j_1}^*, r_{j_1}^*)$ was never signed. Therefore, $\sigma_{j_1}^*$ is a signature on a fresh message, so B wins.
- All of the r_j^* lie in \mathcal{R} , and two of them are equal. Assume without loss of generality that $r_0^* = r_1^* = r_{i_0}$. If $m_0^* = m_1^*$, then we must have $\sigma_0^* \neq \sigma_1^*$, so one of these is a fresh signature. If $m_0^* \neq m_1^*$, then $h_0^* \neq h_1^*$, so one of h_0^* and h_1^* was never signed. Therefore, B also wins.

Since \mathcal{S} is secure, B wins with negligible probability, meaning $(\epsilon - 2/3p)/q_S \ell^2 < \text{negl}$. This is equivalent to $\epsilon - 2/3p < \text{negl}$. Since ϵ is bounded by $1/p$ infinitely often, we have that $1/3p < \text{negl}$ infinitely often, contradicting the fact that p is a polynomial. \square

We note that, for one-time security, $q = 1$, so we only need to rely on the one-time security of \mathcal{S}_c .

A.2 Proof of Theorem 3.19

Proof. We prove security for the stateless case, the stateful case being almost identical. Suppose we have an adversary A that breaks the EUF-qCMA security of \mathcal{S} with non-negligible probability ϵ . We will prove security through a sequence of games.

Game 0. This is the standard attack game, where A makes q quantum queries. For $j = 1, \dots, q$, let \mathbf{a}^j be the vector generated for query j .

Notice that for any \mathbf{b} with $|\mathbf{b}| \leq d - 1$, we only use the secret key $\text{sk}_{\mathbf{b}}$ to sign a single classical message.

Game 1. We now replace $F(k, \cdot)$ with a truly random function. The security of F implies that A still wins with probability negligibly-close to ϵ .

Game 2. Now we remove the random function all together. Instead, we keep a table mapping strings \mathbf{b} to tuples $(\text{sk}_{(\mathbf{b},0)}, \text{sk}_{(\mathbf{b},1)}, \text{pk}_{(\mathbf{b},0)}, \text{pk}_{(\mathbf{b},1)}, \sigma_{\mathbf{b}})$. To answer the j th query, pick a random \mathbf{a}^j . For each $i = 1, \dots, d - 1$,

- Let $\mathbf{b} = \mathbf{a}_{[1,i]}^j$, and look up the tuple for \mathbf{b} .
- If the tuple exists, we do nothing.
- If the tuple does not exist:
 - Sample $(\text{sk}_{(\mathbf{b},b)}, \text{pk}_{(\mathbf{b},b)}) = \text{G}_{ot}(\lambda)$ for $b = 0, 1$.
 - Obtain $\text{sk}_{\mathbf{b}}$ by looking up the tuple for $\mathbf{b}_{[1,i-1]}$.
 - Generate $\sigma_{\mathbf{b}} = \text{Sign}_{ot}(\text{sk}_{\mathbf{b}}, (\text{pk}_{(\mathbf{b},0)}, \text{pk}_{(\mathbf{b},1)}))$
 - Associate \mathbf{b} with the tuple $(\text{sk}_{(\mathbf{b},0)}, \text{sk}_{(\mathbf{b},1)}, \text{pk}_{(\mathbf{b},0)}, \text{pk}_{(\mathbf{b},1)}, \sigma_{\mathbf{b}})$.

Game 3. In this game, we abort if we ever have $\mathbf{a}^j = \mathbf{a}^{j'}$ for $j' \neq j$. There are a total of q different \mathbf{a}^j vectors, and they are drawn from a set of size 2^d . Therefore, the probability of abort is at most $q^2/2^{d+1}$, which is negligible. Therefore, A still wins with probability negligibly close to ϵ .

Notice that in Game 3, since all of the \mathbf{a}^j are distinct, we are only using any particular $\text{sk}_{\mathbf{a}}$ key at most once. The adversary produces $q + 1$ distinct $(m_\ell, \Sigma_\ell, \sigma_\ell)$ pairs. There are two distinct possibilities:

- One or more of the Σ_ℓ is outside all of the $\Sigma_{\mathbf{a}^j}$. In this case, one of the signatures in Σ_ℓ is a forgery for one of the public keys generated in answering the signing queries. We can construct a forger for \mathcal{S}_{ot} by randomly guessing which of the public keys will be forged, plugging the given public key into that key, and randomly generating all of the other keys ourselves. Such an forger will successfully forge with probability only polynomially smaller than the probability Σ_ℓ lies outside of the $\Sigma_{\mathbf{a}^j}$. The assumption that \mathcal{S}_{ot} is secure shows that this probability is therefore negligible.
- Two of the Σ_ℓ are identical. In this case, there is an \mathbf{a} such that we have two forgeries relative to $\text{sk}_{\mathbf{a}}$. We can similarly construct a forger for \mathcal{S}_{ot} by guessing a random \mathbf{a} , and plugging in the given public key as $\text{pk}_{\mathbf{a}}$, and generating the rest of the keys itself. Such a forger will successfully forge with probability only polynomially-smaller than the probability that two of the σ_ℓ are identical. The security of \mathcal{S}_{ot} shows that this probability is also negligible.

Therefore, the probability that A wins in Game 3 is negligible, meaning ϵ is negligible. Hence, \mathcal{S} is secure. □

B Technical Proofs

B.1 Proof of Lemma 2.1

We prove Lemma 2.1, which states that taking a partial measurement obtaining one of k outcomes during a computation only decreases any output's probability by at most a factor of k .

Proof. Let $|\psi\rangle$ be the final state of A , and let $|\psi_y\rangle$ be the final state of A' when the outcome of the partial measurement is y . Let $\Pr[y]$ be the probability that the partial measurement obtains y . It is straightforward to show that $|\psi\rangle = \sum_y \sqrt{\Pr[y]} \alpha_y |\psi_y\rangle$ for some α_y of unit norm. Then we have

$$\Pr[x] = |\langle x|\psi\rangle|^2 = \left| \sum_y \sqrt{\Pr[y]} \alpha_y \langle x|\psi_y\rangle \right|^2 \leq k \sum_y \Pr[y] |\langle x|\psi_y\rangle|^2 = k \Pr'[x]$$

□

B.2 Proof of Lemma 2.5

Recall that we have two sets \mathcal{X} and \mathcal{Y} , and for each $x \in \mathcal{X}$, distributions D_x and D'_x on \mathcal{Y} such that $|D_x - D'_x| \leq \epsilon$ for all x . Let $O : \mathcal{X} \rightarrow \mathcal{Y}$ be a function where, for each x , $O(x)$ is drawn from D_x , and let $O'(x)$ be a function where, for each x , $O'(x)$ is drawn from D'_x . We wish to bound the distinguishing probability of the functions O and O' .

We first suppose that each of the probabilities in each of the distributions D_x and D'_x are rational.

Claim B.1. *If each of the probabilities in D_x and D'_x are rational, then any quantum algorithm making q quantum queries can only distinguish O from O' with probability $\sqrt{8C_0q^3}\epsilon$.*

Before proving this claim, we explain how it proves Lemma 2.5. Fix any quantum algorithm A . The distinguishing probability for any rational collection of distributions D_x and D'_x is bounded by $\sqrt{8C_0q^3 \max_x |D_x - D'_x|}$. But the distinguishing probability of A is a continuous function of the probabilities in the distributions D_x and D'_x , and the pairs of rational distributions are dense in the set of all pairs of distributions. Therefore, the bound of $\sqrt{8C_0q^3 \max_x |D_x - D'_x|}$ applies for all pairs of distributions.

Now we prove the claim:

Proof. Let r be the smallest integer such that each of the probabilities in each of the distributions D_x and D'_x can be represented as a rational number with denominator r . Observe that we can take ϵ to be an integer times $2/r$, say $2s/r$. Let $\mathcal{Z} = [s + r]$. Let E be the uniform distribution on $[r]$ and E' the uniform distribution on $[r] + s = \{s + 1, \dots, s + r\}$. The probabilities in E and E' are the same on $[s] + (r - s) = \{s, \dots, r\}$, and are $1/r$ on $[s]$ and $[s] + r$ respectively. Therefore $|E - E'| = 2s/r$. We now construct functions f_x such that if $z \leftarrow E$, $f_x(y)$ is distributed according to D_x and if $z \leftarrow E'$, $f_x(y)$ is distributed according to D'_x . For each $y \in \mathcal{Y}$, let p/r be the probability under D_x and p'/r the probability under D'_x . Suppose $p \leq p'$. Then we will choose p elements of $[s] + (r - s)$ that have not been chosen before, and let f_x evaluate to y on those elements. We will also choose $p' - p$ elements of $[s] + r$ and let f_x be y on those elements as well. We treat the $p' < p$ case similarly. Then f_x evaluates to y with the desired probabilities, so it remains to show that we never run out of elements. Since $|D_x - D'_x| \leq 2s/r$, we will never run out of elements in $[s] + r$ or $[s]$. If $|D_x - D'_x| < 2s/r$, we will run out of elements in $[s] + (r - s)$. When we run out, however, instead of picking an element in $[s] + (r - s)$, we can pick two elements, one in each of $[s]$ and $[s] + r$, and still have the correct probability.

Now that we can generate $D_x = f_x \circ E$ and $D'_x = f_x \circ E'$, we can generate O and O' differently. Let P be the set of oracles from \mathcal{X} to \mathcal{Z} where each output is drawn according to E , and let P' be the set of oracles where each output is drawn from E' . Then letting $O(x) = f_x(P(x))$ and $O'(x) = f_x(P'(x))$ gives the correct distributions for O and O' . Suppose A distinguishes O from O' with probability σ . Then we can easily construct an algorithm B that distinguishes P and P' with probability σ .

Let ℓ be some integer to be chosen later. We replace P and P' with small-range distributions on ℓ samples of E and E' respectively. Applying Lemma 2.4 twice shows that B must still distinguish P and P' with probability at least $\sigma - 2C_0q^3/\ell$. But now the difference between the distribution P and P' is only ℓ samples of either E or E' , so the distinguishing probability is at most $\ell\epsilon$. Thus $\sigma \leq \ell\epsilon + 2C_0q^3/\ell$ for any ℓ . Setting $\ell = \sqrt{2C_0q^3/\epsilon}$ minimizes this quantity, yielding $\sqrt{8C_0q^3}\epsilon$ as desired. \square

B.3 Proof of Lemma 2.6

Recall that we have sets \mathcal{X} and \mathcal{Y} , and a distribution D on \mathcal{Y} . Let H be a function from \mathcal{X} to \mathcal{Y} where each output is drawn independently according to D . Let F be a function from \mathcal{Y} to some set \mathcal{Z} , and let H_∞ be the minimum over all $z \in \mathcal{Z}$ of the min-entropy of the distribution D on y

conditioned on $F(y) = z$. Let A be a quantum algorithm making q queries to H and any number of queries to $F \circ H$. We wish to show that A can only produce $q + 1$ distinct input/output pairs with probability $(q + 1)/\lceil 2^{H_\infty} \rceil$

We proceed by converting an algorithm violating Lemma 2.6 to an algorithm violating the following lemma proved by Boneh and Zhandry [BZ13]

Lemma B.2 ([BZ13]). *Fix sets \mathcal{X} and \mathcal{Y} , and let H be a random function from \mathcal{X} to \mathcal{Y} . Then any quantum algorithm making q quantum queries can only produce $q + 1$ input/output pairs with probability at most $(q + 1)/|\mathcal{Y}|$.*

First, we need the following technical lemma:

Lemma B.3. *Fix and integer r . Let D be a distribution of a set \mathcal{X} such that $\Pr[x \leftarrow D] < 1/r$ for all x . Then we can construct a distribution D' on injective functions from $[r]$ into \mathcal{X} with the property that $\Pr[x : f \xleftarrow{R} D', i \xleftarrow{R} [r], x \leftarrow f(i)] = \Pr[x : x \xleftarrow{R} D]$ for all y . In other words, we can generate x according to D by drawing a random value i in $[r]$, a random injective function f from D' , and evaluating $f(i)$.*

Proof. Pick an arbitrary ordering of elements in \mathcal{X} . Then there is a one-to-one correspondence between subsets of \mathcal{X} of size r and strictly monotonically increasing functions from $[r]$ to \mathcal{X} . Therefore, it suffices to show how do sample subsets $T \subseteq \mathcal{X}$ of size r such that sampling T and then picking a random element of T simulates the distribution D . We give the algorithm *SampleSubset*, which takes as input a set \mathcal{X} , a distribution D on \mathcal{X} , and an integer r such that where $\Pr[x \leftarrow D] \leq 1/r$, and samples from a distribution of subsets of size r with the desired properties:

We now prove that *SampleSubsets* works as promised. We need to show that D' and D'' are distributions. Since p^* is at most the smallest probability in D , all the probabilities in D' are non-negative. Moreover, by adding up all the probabilities in D' , we see that they sum to 1, so D' is in fact a distribution. This means all the probabilities in D'' are non-negative as well. Using the fact that all elements in \mathcal{F} have probability $1/r$ under D' , we see that the probabilities in D'' also sum to 1, so D'' is also a distribution. The fact that D' is a distribution also shows that $|\mathcal{F}| \leq r$, since otherwise the probabilities would sum to greater than 1.

Next, we explain why the recursive call to *SampleSubset* is valid. That is, that $\Pr[x \leftarrow D''] \leq \frac{1}{r'}$ where $r' = r - |\mathcal{F}|$. For D' , the maximum probability is at most

$$\frac{p_H}{1 - rp^*} \leq \frac{p_H}{1 - (\frac{1}{r} - p_H)} = \frac{r}{r + (r - 1)/p_H} \leq \frac{r}{r + (r - 1)r} = \frac{1}{r}$$

For D'' , the maximum probability is less than $\frac{1/r}{1 - |\mathcal{F}|/r} = \frac{1}{r - |\mathcal{F}|} = \frac{1}{r'}$, as desired. Also, *SampleSubset* is never called with $r' = 0$, since in this case we would have already outputted \mathcal{F} .

Now, we need to show that this sampling algorithm actually terminates. We look at two cases:

- $p^* = p_L$. Let x_L be an element in \mathcal{T} with $\Pr[x_L \leftarrow D] = p_L$. Observe that under D' and hence D'' , x_L has probability 0.
- $p^* = \frac{1}{r} - p_H$. Let x_H be an element with $\Pr[x_H \leftarrow D] = p_H$. Under D' , $\Pr[x_H \leftarrow D'] = \frac{1}{r}$, so x_H is included in \mathcal{F} . Therefore, under D'' , x_H has probability 0.

Algorithm 1 *SampleSubset*(\mathcal{X}, D, r)

If $r = 1$, draw $x \leftarrow D$, output $\{x\}$, and exit.

Otherwise, let p_L be the smallest non-zero probability in D .

Let p_H be the largest probability in D .

Let $p^* \leftarrow \min(p_L, \frac{1}{r} - p_H)$.

Let \mathcal{T} be the set of the r elements in \mathcal{X} with the smallest non-zero probabilities.

With probability rp^* , output \mathcal{T} and exit.

Otherwise, let D' be the distribution where

$$\Pr[x \leftarrow D'] = \begin{cases} \frac{\Pr[x \leftarrow D] - p^*}{1 - rp^*} & \text{if } x \in \mathcal{T} \\ \frac{\Pr[x \leftarrow D]}{1 - rp^*} & \text{otherwise} \end{cases}$$

Let \mathcal{F} be the set of x such that $\Pr[x \leftarrow D'] = \frac{1}{r}$.

If $|\mathcal{F}| = r$, then output \mathcal{F} and exit.

Otherwise, let D'' be the distribution where

$$\Pr[x \leftarrow D''] = \begin{cases} 0 & \text{if } x \in \mathcal{F} \\ \frac{\Pr[x \leftarrow D']}{1 - |\mathcal{F}|/r} & \text{otherwise} \end{cases}$$

Sample T_0 from *SampleSubset*($\mathcal{X}, D'', r - |\mathcal{F}|$)

Output $T_0 \cup \mathcal{F}$.

This means that in each recursive call to *SampleSubset*, the number of x with positive probability decreases by at least 1. Since \mathcal{X} is finite, eventually, the number of x with positive probability will equal r (it cannot be less since all probabilities in D are at most than $1/r$, meaning there are at least r such elements).

It remains to be proven that our sampling algorithm gives the desired distribution. In the case $r = 1$, then we just output sets $\{x\}$ where $x \leftarrow D$, which is correct. Otherwise, with probability rp^* , we output \mathcal{T} . In this case, drawing a random value from \mathcal{T} gives us each element with probability p^* . Since p^* is at most p_L , we have not over-sampled any element. If we do not output \mathcal{T} , we then need to sample subsets to match the distribution D' . If any x has $\Pr[x \leftarrow D'] = \frac{1}{r}$, then x must be in every subset, so we set it aside in the set \mathcal{F} . We then need to draw $r' = r - |\mathcal{F}|$ additional elements not in \mathcal{F} to match the correct distribution. It is straightforward to show that this is achieved by calling *SampleSubset*($\mathcal{X}, D'', r - |\mathcal{F}|$).

□

We are now ready to prove Lemma 2.6:

Proof. Recall that we have an algorithm making q queries to a random oracle H where outputs are drawn from a distribution D , and an arbitrary number of queries to $F \circ H$, and produces $q + 1$ input/output pairs with probability ϵ . Additionally, for all z , $\Pr[y : y \leftarrow D, F(y) = z] \leq p = 2^{-H_\infty}$.

We now generate the two oracles in a different way: Let $O(x)$ be an oracle with the same distribution as $F \circ H$. Fix some arbitrary ordering on \mathcal{Y} . Let $G(z, i)$ output the i th value $y \in \mathcal{Y}$ such that $F(y) = z$. We then generate H as $H(x) = G(O(x), R_O(x))$ where $R_O(x)$ is an oracle

(dependent on O) mapping \mathcal{X} to the positive integers, where for each x , the output distribution is chosen so that $G(O(x), R_O(x))$ has the distribution matching D , conditioned on $F(y) = O(x)$.

Thus H and O have the correct distributions, and A is a quantum algorithm making q queries to H and many queries to O , and outputting $q + 1$ pairs x_i, y_i where $y_i = H(x_i) = G(O(x_i), R_O(x_i))$ with probability ϵ . Notice from y_i , we can recover $R_O(x_i)$ by just counting how many $y \leq y_i$ have $F(y) = F(y_i)$.

Next, we change how $R_O(x)$ is drawn. Notice that, by the min-entropy assumption, $R_O(x) = j$ with probability at most p for any j . Let $r = \lfloor 1/p \rfloor$. Then $R_O(x) = j$ with probability at most $1/r$. Therefore, according to Lemma B.3, we can draw $R_O(x)$ by first drawing an injective function $S_{O,x} : [r] \rightarrow \mathbb{Z}$ from an appropriate distribution (where the distribution depends on O and x), picking a random $k \in [r]$, and outputting $S_{O,x}(k)$. Doing this for all x , we can generate H as $H(x) = G(O(x), S_{O,x}(P(x)))$ where P is a random function from \mathcal{X} to $[r]$.

From $R_O(x_i) = S_{O,x_i}(P(x_i))$, we can recover $P(x_i)$ since S_{O,x_i} is injective. Notice that to answer all the queries requires only q queries to P , and we can produce $q + 1$ input output pairs of P with probability ϵ . Lemma B.2 gives us the bound $\epsilon \leq (q + 1)/r$, which completes the proof. \square