Secure Signatures and Chosen Ciphertext Security in a Quantum Computing World

Dan Boneh and 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 ubiquitous quantum computing 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, quantum security.

1 Introduction

Recent progress in building quantum computers [IBM12] gives hope for their eventual feasibility. Consequently, there is a growing need for quantum-secure 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

R. Canetti and J.A. Garay (Eds.): CRYPTO 2013, Part II, LNCS 8043, pp. 361–379, 2013.

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 [BZ13a] 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 [BZ13a] gives the attacker the signatures on all messages in a quantum superposition. In more detail, a quantum chosen message query is the transformation

$$\sum_m \psi_m \ |m\rangle \qquad \longrightarrow \qquad \sum_m \psi_m \ |m, \ S(\mathsf{sk}, m)\rangle$$

where $S(\mathsf{sk}, x)$ 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 how to transform any signature that is existentially unforgeable under a *classical random* message into a signature scheme that is 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. These results 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 is non-optimized, and can potentially be made more efficient using standard combinatorial techniques. Unfortunately, we cannot prove quantum-security of an optimized Lamport signature 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 challenge ciphertexts. In the quantum setting, the decryption oracle accepts a superposition of ciphertexts and returns a superposition of their decryptions:

$$\sum_m \psi_c \, |c\rangle \qquad \longrightarrow \qquad \sum_c \psi_c \, |c, \; D(\mathsf{sk}, c)\rangle \; .$$

One might also try to allow quantum access to the encryption oracle; however, we show that the resulting concept is unsatisfiable. 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. This construction is the quantum analogue of the CHK transformation from identity-based encryption to public-key chosen ciphertext security [BCHK04].

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 all computers are quantum, the last step in a signature or decryption procedure is to sample the final quantum state. This ensures that the results are always classical, thereby preventing quantum superposition attacks. Security in this case relies on a physical hardware assumption, namely that the final "classicalization" step is implemented correctly and cannot be circumvented by a quantum adversary. In contrast, using systems that are inherently secure against superposition attacks frees the hardware designer from worrying about the security of the classicalization step.

As further motivation, we note that our results are the tip of a large emerging area with many open questions. For any cryptographic primitive modeled as an interactive game, one can ask how to design primitives that remain secure when the interaction between the adversary and its given oracles is quantum. For example, can we design quantum-secure threshold signatures and group signatures? Can we construct a quantum-secure PRF for a large domain from a quantum-secure PRF for a small domain? In particular, do the CBC-MAC or NMAC constructions give quantum-secure PRFs?

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 setting [BS08, BHK+11]. Damgård 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 composability (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 and do not consider quantum superposition attacks.

2 Preliminaries: Background and Techniques

We will let $[n]$ denote the set $\{1, ..., n\}$. Functions will be denoted by capital letters (such as F), and sets by capital script letters (such as $\mathcal{X}$). We will let $x \xleftarrow{R} D$ for some distribution D denote drawing x according to D, and $x \xleftarrow{R} \mathcal{X}$

for some set $\mathcal{X}$ denote drawing a random element from $\mathcal{X}$. Given a function $F : \mathcal{X} \to \mathcal{Y}$ and a subset $\mathcal{S} \subseteq \mathcal{X}$, the restriction of F to $\mathcal{S}$ is the function $F_{\mathcal{S}} : \mathcal{S} \to \mathcal{Y}$ 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. A non-negative function $f(n)$ is negligible if, for any c, $f(n) < 1/n^c$ for all sufficiently large n. If a function $g(n)$ can be written as $h(n) \pm f(n)$ where $f(n)$ is negligible, we write $g(n) = h(n) \pm \mathsf{negl}$.

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 and 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 $B = \{|b_0\rangle, ..., |b_{d-1}\rangle\}$ for $\mathcal{H}$, a measurement of $|\psi\rangle$ in the basis B results in the value i with probability $|\langle b_i | \psi \rangle|^2$, and the quantum state collapses to the basis vector $|b_i\rangle$. If $|\psi\rangle$ is actually a state in a joint system $\mathcal{H} \otimes \mathcal{H}'$, then $|\psi\rangle$ can be written as

$$|\psi\rangle = \sum_{i=0}^{d-1} \alpha_i \, |b_i\rangle |\psi_i'\rangle$$

for some complex values α_i and states $|\psi_i'\rangle$ over $\mathcal{H}'$. Then, the measurement over $\mathcal{H}$ obtains the value i with probability $|\alpha_i|^2$ and in this case the resulting quantum state is $|b_i\rangle|\psi_i'\rangle$.

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 $\mathbf{U}_1, ..., \mathbf{U}_n$ 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

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

The final state is then measured, obtaining the tuple (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 $O : \mathcal{X} \to \mathcal{Y}$ by a unitary transformation $\mathbf{O}$ where

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

where $+ : \mathcal{X} \times \mathcal{X} \to \mathcal{X}$ is some group operation on $\mathcal{X}$. Suppose we have a quantum algorithm that makes quantum queries to oracles $O_1, ..., O_q$. Let $|\psi_0\rangle$ be the input state of the algorithm, and let $\mathbf{U}_0, ..., \mathbf{U}_q$ be the unitary transformations applied between queries. Note that the transformations $\mathbf{U}_i$ are themselves possibly the products of many simpler unitary transformations. The final state of the algorithm will be

$$\mathbf{U}_q\mathbf{O}_q...\mathbf{U}_1\mathbf{O}_1\mathbf{U}_0|\psi_0\rangle$$

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 $\mathbf{O}_i$. We call a quantum oracle algorithm efficient if the number of queries q is a polyomial, and each of the transformations $\mathbf{U}_i$ between queries can be written as the product polynomially many unitary transformations from some fixed basis set.

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 the full version [BZ13b]. The first concerns partial measurements, and will be used extensively throughout the paper:

Lemma 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 on the state of A 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 ([Zha12a]). *Let H be an oracle drawn from a $2q$-wise independent distribution. Then the advantage any quantum algorithm making at most q queries to H has in distinguishing H from a truly random function is identically 0.*

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

Lemma 3 ([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 $H : \mathcal{X} \to \mathcal{Y}$, the following two cases are indistinguishable, except with probability less than C_0q^3/ℓ:*

- *$H(x) = y_x$ 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.*

3 Quantum-Secure Signatures

Our goal is to construct signatures that are resistant to a *quantum* chosen message attack, where the adversary submits quantum superpositions of messages and receives the corresponding superpositions of signatures in return. First, we need a suitable definition of what a signature scheme is in our setting, and what it means for such a scheme to be secure. Correctness for a stateless signature scheme is identical to the classical setting: any signature produced by the signing algorithm must verify. There is some subtlety, however, for stateful signature schemes. If the state of the signing algorithm depends on the messages signed, and if the adversary mounts a quantum chosen message attack, the signing algorithm and adversary will become entangled. To keep the state of the signing algorithm classical and unentangled with the adversary, we therefore restrict the state to be independent of the messages signed so far. We note that many stateful signature schemes, such as stateful Merkle signatures, satisfy this requirement. We arrive at the following definition:

Definition 2. *A signature scheme $\mathcal{S}$ is a tuple of efficient classical algorithms* $(\mathsf{G}, \mathsf{Sign}, \mathsf{Ver})$ *where*

- $\mathsf{G}(\lambda)$ *generates a private/public key pair* $(\mathsf{sk}, \mathsf{pk})$.
- $\mathsf{Sign}(\mathsf{sk}, m, \mathit{state})$ *outputs a signature* σ *and new state* state'. *If the output* state *is ever non-empty, we say that algorithm* Sign *is stateful and we require that the state does not depend in any way on the messages that have been signed so far. If the output* state *is always empty, we say that* Sign *is stateless and we drop the* state *variables altogether.*
- $\mathsf{Ver}(\mathsf{pk}, m, \sigma)$ *either accepts or rejects. We require that valid signatures are always accepted, that is if* σ *is the output of* $\mathsf{Sign}(\mathsf{sk}, m, \mathit{state})$ *then* $\mathsf{Ver}(\mathsf{pk}, m, \sigma)$ *accepts.*

For security, we use a notion similar to that for message authentication codes defined by Boneh and Zhandry [BZ13a]. There are two issues in defining security under a quantum chosen message attack:

- **Randomness.** When using a randomized signature scheme, there are several choices for how the randomness is used. One option is to choose a single randomness value for each chosen message query, and sign every message in the superposition with that randomness. Another approach is to choose fresh randomness for each message in the superposition. Using a single randomness value for each query is much simpler for implementers, and we therefore design signature schemes secure in this setting.
  Fortunately, there is a simple transformation that converts a scheme requiring independent randomness for every message into a scheme that is secure when a single randomness value is used for an entire query: when signing, choose a fresh random key k for a quantum pseudorandom function (QPRF). This will be the single per-query randomness value. To sign a superposition of messages, sign each message m in the superposition using randomness

obtained by applying the QPRF to m using the key k. From the adversary's point of view, this is indistinguishable from choosing independent randomness for each message. Using Lemma 2, we can replace the QPRF with a function drawn from a pairwise independent function family, which is far more efficient than using a QPRF. Hence, requiring global randomness per query does not complicate the signature scheme much, but greatly simplifies its implementation.

– **Forgeries.** Each quantum chosen message query can be a superposition of every message in the message space. Sampling the returned superposition will result in a single message/signature pair for a random message. Therefore, the classical notion of existential forgery being a signature on a *new* message is ill-defined when we allow quantum access. Instead, for security we require that the adversary cannot produce $q+1$ valid message/signature pairs with q quantum chosen message queries. Security definitions in this style were previously used in the context of blind signatures [PS96].

We arrive at the following definition of security:

Definition 3 (Quantum Security). *A signature scheme $\mathcal{S} = (\mathsf{G}, \mathsf{Sign}, \mathsf{Ver})$ 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 $(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{G}(\lambda)$, 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:*

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

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.* □

In this paper, we will also be using several weaker notions of security. The first is for a classical chosen message attack:

Definition 4. *$\mathcal{S}$ is existentially unforgeable under a* classical *random message attack (EUF-CMA secure) if every signing query is measured before signing, so that only a single classical message is signed per query.*

Next, we define random message security:

Definition 5. *$\mathcal{S}$ 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 uniform random messages at the beginning of the game.*

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

Definition 6. *$\mathcal{S}$ is universally unforgeable under a random message attack (UUF-RMA secure) if, along with receiving q message/signature pairs for random messages, the adversary receives n additional random messages, and all of the $q+1$ messages for which a signature is forged must be among the $q+n$ messages received.*

All of the above security definitions also have weak variants, where in addition to requiring that message/signature forgery pairs be distinct, we also require that the messages themselves 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 unbounded q notion as many-time security.

Separation from Classical Security. In the full version [BZ13b], we present a signature scheme that is secure under classical queries, but completely insecure once an adversary can make quantum queries.

The idea is to augment a classically secure scheme by choosing a random secret prime p and storing p in the secret signing key. We modify the signature scheme so that the signature on the message $m = p$ includes the entire secret key. As long as the adversary does not learn p, she should not be able to learn the secret key. Following ideas from Zhandry [Zha12b], we also add some auxiliary information to the signatures such that, under classical queries, p is hidden, but a single quantum query suffices to recover p. Since classically, signatures can be built from one-way functions, we immediately get the following theorem:

Theorem 1. *Assuming the existence of one-way functions, there exists a signature scheme $\mathcal{S}$ that is existentially unforgeable under a classical chosen message attack, but is totally broken under a quantum chosen message attack.*

3.1 Quantum-Secure Signatures from Classically-Secure Signatures

Now we move to actually building signature schemes that are secure against quantum chosen message attacks. 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 7. *A chameleon hash function $\mathcal{H}$ is a tuple of efficient algorithms* $(\mathsf{G}, \mathsf{H}, \mathsf{Inv}, \mathsf{Sample})$ *where:*

- $\mathsf{G}(\lambda)$ *generates a secret/public key pair* $(\mathsf{sk}, \mathsf{pk})$.
- $\mathsf{H}(\mathsf{pk}, m, r)$ *maps messages to some space $\mathcal{Y}$*

- $\mathsf{Sample}(\lambda)$ *samples* r *from some distribution such that, for every* pk *and* m, $\mathsf{H}(\mathsf{pk}, m, r)$ *is uniformly distributed.*
- $\mathsf{Inv}(\mathsf{sk}, h, m)$ *produces an* r *such that* $\mathsf{H}(\mathsf{pk}, m, r) = h$*, and* r *is distributed negligibly-close to* $\mathsf{Sample}(\lambda)$ *conditioned on* $\mathsf{H}(\mathsf{pk}, m, r) = h$

We say that a chameleon hash function is collision resistant if no efficient quantum algorithm, given only pk, can find collisions in $\mathsf{H}(\mathsf{pk}, \cdot, \cdot)$. 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. The idea behind our construction is to first hash the message with the chameleon hash function and then sign the hash. In order to be secure against quantum queries, care has to be taken in how the randomness for the hash and the signature scheme is generated. In what follows, for any randomized algorithm A, we let $A(x; r)$ denote running A on input x with randomness r.

Construction 2. *Let* $\mathcal{H} = (\mathsf{G}_H, \mathsf{H}, \mathsf{Inv}, \mathsf{Sample})$ *be a chameleon hash function, and* $\mathcal{S}_c = (\mathsf{G}_c, \mathsf{Sign}_c, \mathsf{Ver}_c)$ *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* $\mathcal{S} = (\mathsf{G}, \mathsf{Sign}, \mathsf{Ver})$ *where:*

$$\begin{aligned}
\mathsf{G}(\lambda) : {} & (\mathsf{sk}_H, \mathsf{pk}_H) \xleftarrow{R} \mathsf{G}_H(\lambda),\ (\mathsf{sk}_c, \mathsf{pk}_c) \xleftarrow{R} \mathsf{G}_c(\lambda) \\
& \textit{output } \mathsf{sk} = (\mathsf{pk}_H, \mathsf{sk}_c),\ \mathsf{pk} = (\mathsf{pk}_H, \mathsf{pk}_c) \\
\mathsf{Sign}((\mathsf{pk}_H, \mathsf{sk}_c), m) : {} & Q \xleftarrow{R} \mathcal{Q}, R \xleftarrow{R} \mathcal{R} \\
& r \leftarrow \mathsf{Sample}(\lambda; R(m)),\ s \leftarrow Q(m),\ h \leftarrow \mathsf{H}(\mathsf{pk}_H, m, r) \\
& \sigma \leftarrow \mathsf{Sign}(\mathsf{pk}_c, h; s),\ \textit{output } (r, \sigma) \\
\mathsf{Ver}((\mathsf{pk}_H, \mathsf{pk}_c), m, (r, \sigma)) : {} & h \leftarrow \mathsf{H}(\mathsf{pk}_H, m, r),\ \textit{output } \mathsf{Ver}(\mathsf{pk}_c, h, \sigma)
\end{aligned}$$

We note that the chameleon secret key is not used in Construction 2, 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. *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 2 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 shows that we can take a classically EUF-RMA secure signature scheme, combine it with a a chameleon hash, and obtain a quantum-secure signature scheme. In particular, the following constructions will be quantum secure, assuming SIS is hard:

- A slight modification to the signature scheme of Cash et al. [CHKP10], which combines their chameleon hash function with an EUF-RMA secure signature

scheme. The only difference in their scheme is that the values r and s are sampled directly, rather than setting them to be the outputs of pairwise independent functions.

– A modification of the 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:

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: use the chameleon secret key sk_H to compute a randomness r such that $\mathsf{H}(\mathsf{pk}_H, m, r) = h$, and output the signature (r, σ). Thus, we can respond to a classical chosen message attack, given only signatures on random messages.

If the adversary issues a *quantum* chosen message query, we need to sign each of the exponentially many messages in the query superposition. Therefore, using the above technique directly would require signing an exponential number of random hashes. Instead, we use small-range distributions and Lemma 3 to reduce the number of signed hashes required 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 can, however, argue that two of the forgeries must, in some sense, correspond to the same query. If we knew which query, we could perform a measurement, observing which of the (polynomially many) random hashes were signed. Lemma 1 shows that the adversary's advantage is reduced by only a polynomial factor. For this query, we now only sign a single random hash, but the adversary produces two forgeries. Therefore, one of these forgeries must be a forgery for $\mathcal{S}_c$. Of course, we cannot tell ahead of time which query to measure, so we just pick the query at random, and succeed with probability $1/q$.

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 2 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 $r_m^{(i)} = \mathsf{Sample}(\lambda; R^{(i)}(m))$ and $s_m^{(i)} = Q^{(i)}(m)$.
– Compute $h_m^{(i)} = \mathsf{H}(\mathsf{pk}_H, m, r_m^{(i)})$
– 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 $p = p(\lambda)$ 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. According to Lemma 2, the view of the adversary is unchanged. 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. Additionally, draw uniform 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, and we show in the full version [BZ13b] that this implies Games 1 and 2 are indistinguishable. Therefore, the success probability is at least $\epsilon - \mathsf{negl}$.

Game 3. Let $\ell = 2C_0qp$ where C_0 is the constant from Lemma 3. At the beginning of the game, for $i = 1, ..., q$ and $j = 1, ..., \ell$, 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 3 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 $\mathcal{H}^{(i)} = \{\hat{h}_j^{(i)}\}$ be the set of $\hat{h}$ values used to answer query 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. This means that one of the h_k^* was never signed, or one of them was signed once, but two signatures were produced for it. In either case, it is straightforward to construct a forger B_0 for $\mathcal{S}_c$ that wins in this case. Since $\mathcal{S}_c$ is secure, this event only happens with negligible probability.
- All of the h_k^* values are 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 $\epsilon - \mathsf{negl} - 1/2p$.

Game 5. Now we guess a random query i^* and 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 $\epsilon/q - \mathsf{negl} - 1/2pq$. 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 Lemma 1 says the adversary's success probability is still at least $\epsilon/q\ell - \mathsf{negl} - 1/2pq\ell$. Notice now that for query i^*, the challenger only needs to sign $\hat{h}_{j^*}^{(i^*)}$, and therefore, one of the $h_b^* = \hat{h}_j^{(i^*)}$ values was never signed.

Game 7. Now guess at the beginning of the game the value of j^*, and at the end, check that the guess was correct. The adversary still wins with probability $\epsilon/q\ell^2 - \mathsf{negl} - 1/2pq\ell^2$.
If the adversary wins in Game 7, it produced two signatures on $\hat{h}^{(i^*)}$ values, while only one of them was signed. It is straightforward to construct a forger B_1 for $\mathcal{S}_c$ that wins in this case. B_1 has success probability $\epsilon/q\ell^2 - \mathsf{negl} - 1/2pq\ell^2$, and the security of $\mathcal{S}_c$ implies that this quantity is negligible. Thus $\epsilon - 1/2p$ is negligible. Since $\epsilon > 1/p$ infinitely often, we then have $1/2p < \mathsf{negl}$ infinitely often, a contradiction. Therefore, ϵ is negligible. □

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 present a simple generic conversion from any classical signature scheme to a scheme secure against quantum chosen message attacks in the quantum random oracle model.

Recall that when a random oracle scheme is implemented in the real-world, the random oracle is replaced by a concrete hash function H, thereby enabling a quantum adversary to evaluate H on a superposition of inputs. Therefore, security proofs in the random oracle model must allow all parties, including the adversary, to issue *quantum* queries to H. This model is called the *quantum* random oracle model [BDF+11] and is the one we use here.

Our construction is quite 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. This construction is very appealing since messages are often hashed anyway before signing. The results in this section then show that only minor modifications to existing schemes are necessary to make them quantum immune.

Construction 4. *Let $\mathcal{S}_c = (\mathsf{G}_c, \mathsf{Sign}_c, \mathsf{Ver}_c)$ 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 $\mathcal{S} =$*

$(\mathsf{G}, \mathsf{Sign}, \mathsf{Ver})$ *where:*

$$\mathsf{G}(\lambda) \;= \mathsf{G}_c(\lambda)$$

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

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

We note that Construction 4 is similar to Construction 2: instead of the chameleon hash $\mathsf{H}(\mathsf{pk}, \cdot, \cdot)$ 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 4, assuming only a very weak form of security for $\mathcal{S}_c$, namely, universal unforgeability under a random message attack (UUF-RMA security):

Theorem 5. *If $\mathcal{S}_c$ is strongly (resp. weakly) UUF-RMA secure, then $\mathcal{S}$ in Construction 4 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.*

We prove Theorem 5 in the full version [BZ13b]. Given that Construction 4 is similar to Construction 2, the security proofs are similar. Now, 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, such as pre-image sampleable functions (PSFs) defined by Gentry et al. [GPV08]. Roughly, PSFs are many-to-one functions F such that, with the secret key, a random pre-image can be sampled. For security, we require that without the secret key, the function is one-way and collision resistant. If we sign a message m by sampling a random pre-image of m, and verify a signature σ by checking that $F(\sigma) = m$, then one-wayness plus collision resistance implies strong UUF-RMA security.

Corollary 1. *If* PSF *is a collision resistant and one-way PSF, then Construction 4 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. In the full version [BZ13b], we also show that the basic GPV signature scheme is secure in the quantum random oracle model, though the proof is very different.

Next, it is straightforward to show that any adversary A breaking the universal unforgeability of $\mathcal{S}_c$ by mounting a random message attack can easily be transformed into an adversary B breaking Construction 4 under a *classical* chosen message attack in the *classical* random oracle model. Together with Theorem 5, we get the following:

Corollary 2. *If $\mathcal{S}$ in Construction 4 is weakly (resp. strongly) existentially unforgeable under a* classical *chosen message attack performed by a quantum adversary, then it is also weakly (resp. strongly) exististentially unforgeable under a* quantum *chosen message attack.*

Therefore, if a scheme matches the form of Construction 4, it is only necessary to prove classical security.

3.3 Signatures from Generic Assumptions

We briefly explain how to construct signatures from generic assumptions. We first construct one-time signatures from one-way functions using the basic Lamport construction [Lam79]. In the classical setting, the next step would be to use target collision resistance to expand the message space. Unfortunately, target collision resistance ceases to make sense in the quantum setting, so we resort to collision resistance to expand the message space. Finally, we plug these one-time signatures into the Merkle signature scheme [Mer87]. The end result is a signature scheme whose quantum security relies only on the existence of collision-resistant functions. The following is proved in the full version [BZ13b]:

Theorem 6. *If there exists a collision-resistant hash function, then there exists a strongly EUF-qCMA secure signature scheme.*

4 Quantum-Secure 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. At a high level, our notion of security allows quantum encryption and decryption queries, but requires challenge queries to be *classical*:

Definition 8. *A symmetric key encryption scheme $\mathcal{E} = (\mathsf{Enc}, \mathsf{Dec})$ 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 $1/2 + \mathsf{negl}$:*

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 $c^* = \mathsf{Enc}(k, m_b)$. 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:*

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

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

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

where

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

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

In the above definition, we need to define the operation $m \oplus \perp$. Since the query responses will XOR $\perp$ with different messages, we need a convention that makes this operation reversible. Taking $\perp$ to be some bit string that lies outside of the message space and $\perp \oplus m$ to be the bitwise XOR will suffice.

Note that we implicitly assume that the decryption algorithm is deterministic. This will be true of our encryption schemes. We note that this is not a limiting assumption since we can make the decryption algorithm deterministic by deriving the randomness for decryption from a PRF applied to the ciphertext. Also, as in the classical case, a simple hybrid argument shows that the above definition is equivalent to the case where the number of encryption queries is limited to one. Lastly, it is straightforward to modify the above definition for public key encryption schemes.

Quantum Challenge Queries. One might hope to enhance Definition 8 by making the security game entirely quantum, where challenge queries are quantum as well. This leads to several difficulties. First, with quantum challenge queries, it is no longer possible to record the challenge ciphertext. This makes it difficult to check that the adversary only asks decryption queries on ciphertexts other than the challenge ciphertexts. The second difficulty is more serious: allowing quantum challenge queries results in definitions of security that are unachievable, even if we disallow decryption queries. In the full version [BZ13b], we show several attempts at defining security with quantum challenges, and show that each of these definitions is insecure.

Separation from Classical Security. Similar to the case for signatures, quantum chosen ciphertext queries give the adversary more power than classical queries. The following is proved in the full version [BZ13b]:

Theorem 7. *If there exists a symmetric (resp. public) key encryption scheme $\mathcal{E}$ that is secure against a* classical *chosen ciphertext attack, then there is a symetric (resp. public) key encryption scheme $\mathcal{E}'$ that is secure under a classical chosen ciphertext attack, but totally insecure under a* quantum *chosen ciphertext attack.*

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 quantum chosen *plaintext* secure encryption scheme and any EUF-qCMA secure MAC, would be quantum chosen *ciphertext* 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 8. *Let F and G be pseudorandom functions. We construct the following encryption scheme $\mathcal{E} = (\mathsf{Enc}, \mathsf{Dec})$ where:*

$$\begin{aligned}\mathsf{Enc}((k_1,k_2),m) : & \; r \xleftarrow{R} \{0,1\}^\lambda \\ & \; c_1 \leftarrow F(k_1,r) \oplus m,\ c_2 \leftarrow G(k_2,(r,m)) \\ & \; \textit{output } (r,c_1,c_2)\end{aligned}$$

$$\begin{aligned}\mathsf{Dec}((k_1,k_2),(r,c_1,c_2)) : & \; m \leftarrow c_1 \oplus F(k_1,r),\ c_2' \leftarrow G(k_2,(r,m)) \\ & \; \textit{if } c_2 \neq c_2',\ \textit{output } \perp \\ & \; \textit{otherwise, output } m\end{aligned}$$

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 9. *If F and G are quantum-secure pseudrandom functions, then $\mathcal{E}$ in Construction 8 is qCCA-secure.*

Theorem 9 is proved in the full version [BZ13b]. As demonstrated by Zhandry [Zha12b], quantum-secure pseudorandom functions can be built from any one-way function. Therefore, Theorem 9 shows that quantum chosen ciphertext security can be obtained from the minimal assumption that one-way functions exist.

4.2 Public-key CCA Security

In the full version [BZ13b], we construct CCA-secure signatures in the public-key setting. We follow the generic transformation from identity-based encryption (IBE) to CCA security due to Boneh et al.[BCHK04], which uses a selectively secure IBE scheme and a strong one-time signature scheme. The one-time signature scheme only needs to be classically secure, and can hence be built from

any one-way function. In contrast, we need the IBE scheme to be secure against *quantum* chosen identity queries. We observe that the IBE scheme of Agrawal, Boneh, and Boyen [ABB10] meets this notion of security, assuming the hardness of the Learning With Errors (LWE) problem. We obtain the following:

Theorem 10. *If the* LWE *problem is hard for quantum computers, then there exists a public-key encryption scheme that is IND-qCCA secure.*

5 Conclusion and Open Problems

We defined the notions of a quantum chosen message attack for signatures and quantum chosen ciphertext attack for encryption. We gave the first constructions of signatures and encryption schemes meeting these strong notions of security. For signatures, we presented two simpler compilers that transform classically secure schemes into quantum-secure schemes. We also showed that signatures can be built from any collision resistant hash function. For encryption, we presented both a symmetric-key and a public-key construction. There are many directions for future work. First, can we base quantum security for signatures on the minimal assumption of one-way functions? Also, it may be possible to mount quantum superposition attacks against many cryptographic primitives. For example, can we build identification protocols or functional encryption that remain secure in the presence of such attacks?

Acknowledgments. We thank Luca Trevisan and Amit Sahai for helpful conversations about this work. This work was supported by NSF, DARPA, the Air Force Office of Scientific Research (AFO SR) under a MURI award, Samsung, and a Google Faculty Research Award. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA or the U.S. Government.

References

[ABB10] Agrawal, S., Boneh, D., Boyen, X.: Efficient lattice (H)IBE in the standard model. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 553–572. Springer, Heidelberg (2010)

[BCHK04] Canetti, R., Halevi, S., Katz, J.: Chosen-ciphertext security from identity-based encryption. In: Cachin, C., Camenisch, J.L. (eds.) EUROCRYPT 2004. LNCS, vol. 3027, pp. 207–222. Springer, Heidelberg (2004)

[BDF+11] Boneh, D., Dagdelen, Ö., Fischlin, M., Lehmann, A., Schaffner, C., Zhandry, M.: Random Oracles in a Quantum World. In: Lee, D.H., Wang, X. (eds.) ASIACRYPT 2011. LNCS, vol. 7073, pp. 41–69. Springer, Heidelberg (2011)

[BHK+11] Brassard, G., Høyer, P., Kalach, K., Kaplan, M., Laplante, S., Salvail, L.: Merkle Puzzles in a Quantum World. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 391–410. Springer, Heidelberg (2011)

[BS08] Brassard, G., Salvail, L.: Quantum Merkle Puzzles. In: Second International Conference on Quantum, Nano and Micro Technologies (ICQNM 2008), pp. 76–79 (February 2008)

[BZ13a] Boneh, D., Zhandry, M.: Quantum-secure message authentication codes. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 592–608. Springer, Heidelberg (2013), Full version available at the Electronic Colloquium on Computational Complexity: http://eccc.hpi-web.de/report/2012/136

[BZ13b] Boneh, D., Zhandry, M.: Secure signatures and chosen ciphertext security in a quantum computing world. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 361–379. Springer, Heidelberg (2013), Full version available at the Cryptology ePrint Archives (2013), http://eprint.iacr.org/2013/088

[Can01] Canetti, R.: Universally composable security: A new paradigm for cryptographic protocols. In: Proceedings of FOCS. IEEE (2001)

[CHKP10] Cash, D., Hofheinz, D., Kiltz, E., Peikert, C.: Bonsai Trees, or How to Delegate a Lattice Basis. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 523–552. Springer, Heidelberg (2010)

[DFNS11] Damgård, I., Funder, J., Nielsen, J.B., Salvail, L.: Superposition attacks on cryptographic protocols. CoRR, abs/1108.6313 (2011)

[GPV08] Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for Hard Lattices and New Cryptographic Constructions. In: Proceedings of the 40th Annual ACM symposium on Theory of computing (STOC), p. 197 (2008)

[HSS11] Hallgren, S., Smith, A., Song, F.: Classical cryptographic protocols in a quantum world. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 411–428. Springer, Heidelberg (2011)

[IBM12] IBM Research. IBM research advances device performance for quantum computing (February 2012), http://www-03.ibm.com/press/us/en/pressrelease/36901.wss

[KR00] Krawczyk, H., Rabin, T.: Chameleon hashing and signatures. In: Proc. of NDSS, pp. 1–22 (2000)

[Lam79] Lamport, L.: Constructing digital signatures from a one-way function. Technical Report SRI-CSL-98 (1979)

[Mer87] Merkle, R.C.: A Digital Signature Based on a Conventional Encryption Function. In: Pomerance, C. (ed.) CRYPTO 1987. LNCS, vol. 293, pp. 369–378. Springer, Heidelberg (1988)

[PS96] Pointcheval, D., Stern, J.: Provably secure blind signature schemes. In: Kim, K.-C., Matsumoto, T. (eds.) ASIACRYPT 1996. LNCS, vol. 1163, pp. 1–12. Springer, Heidelberg (1996)

[Unr10] Unruh, D.: Universally Composable Quantum Multi-Party Computation. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 486–505. Springer, Heidelberg (2010)

[Zha12a] Zhandry, M.: Secure identity-based encryption in the quantum random oracle model. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 758–775. Springer, Heidelberg (2012), Full version available at the Cryptology ePrint Archives: http://eprint.iacr.org/2012/076/

[Zha12b] Zhandry, M.: 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/