

Secure Identity-Based Encryption in the Quantum Random Oracle Model*

Mark Zhandry

Stanford University, USA

Abstract. We give the first proof of security for an identity-based encryption scheme in the *quantum* random oracle model. This is the first proof of security for *any* scheme in this model that requires no additional assumptions. Our techniques are quite general and we use them to obtain security proofs for two random oracle hierarchical identity-based encryption schemes and a random oracle signature scheme, all of which have previously resisted quantum security proofs, even using additional assumptions. We also explain how to remove the extra assumptions from prior quantum random oracle model proofs. We accomplish these results by developing new tools for arguing that quantum algorithms cannot distinguish between two oracle distributions. Using a particular class of oracle distributions, so called *semi-constant* distributions, we argue that the aforementioned cryptosystems are secure against quantum adversaries.

Keywords: Quantum, Random Oracle, IBE, Signatures.

1 Introduction

While quantum computation is not yet viable, Shor [Sho97] showed that when fully realized, quantum computers will break most of the cryptosystems used today, namely those based on the difficulty factoring and the discrete log problem. This has sparked the field of post-quantum cryptography, the search for classical systems secure against quantum adversaries. To be secure in this setting, a system must have an underlying difficult problem for quantum computers, as well as a security reduction showing how to solve this problem using a quantum adversary that breaks the system. Problems based on lattices have so far resisted quantum attacks, and there is a considerable amount of literature on lattice-based cryptosystems. However, random oracle security proofs for these constructions typically only consider classical adversaries, thus failing to show security in the quantum world.

The random oracle model [BR93] is one area where many classical proofs lack quantum equivalents. This model is of interest because random oracle schemes tend to be more efficient than their standard model counterparts. Consequently, the most efficient lattice-based schemes are often constructed in the random oracle model. For example, Gentry, Peikert, and Vaikuntanathan [GPV08]

* Full version available at <http://eprint.iacr.org/2012/076/>

show how to construct signatures and identity-based encryption (IBE). Cash et al. [CHKP10] and Agrawal, Boneh, and Boyen [ABB10] give hierarchical IBE schemes both in the standard model and the random oracle model, with the random oracle constructions being more efficient than the corresponding standard model schemes. Gordon, Katz, and Vaikuntanathan [GKV10] give a group signature scheme, while Boneh and Freeman [BF11] demonstrate a homomorphic signature scheme, all in the random oracle model.

All of these papers prove security against adversaries with classical access to the random oracle. However, when the scheme is instantiated, the random oracle is replaced with a hash function H , which a quantum adversary may evaluate on a quantum superposition of inputs. To model this ability, it is necessary to allow a quantum adversary to make quantum queries to the random oracle. We call this model the quantum random oracle model.

Proving security in the quantum random oracle model presents many challenges. Among them is the difficulty of efficiently simulating the random oracle. In the classical case, the random oracle is simulated on the fly, only generating randomness as needed. In the quantum setting, the adversary can query the oracle on an exponential superposition of inputs. Therefore, to make even the first query look random to the adversary, it would seem that we would need exponential randomness.

In addition to this difficulty, there are classical random oracle techniques that do not make sense if the adversary has quantum access to the random oracle. One such technique is that of Bellare and Rogaway [BR93] for proving the security of the Full Domain Hash signature scheme. In this technique, the reduction algorithm is given a challenge c to solve, and randomly guesses which oracle query the adversary will use to break the scheme. The algorithm embeds c into the response for this query, and if both the guess is correct and the adversary breaks the scheme, then the algorithm will be able to solve c .

In the quantum setting, this argument no longer applies because each oracle query might be over a superposition of exponentially many inputs. We could choose a random query and plug c into all outputs for that query, but this will not look like a random oracle to the adversary. Alternatively, we could choose one oracle input and plug c into the corresponding output for all queries. However, our chance of guessing correctly will then be exponentially small.

1.1 Our Contributions

We resolve some of the issues outlined above by giving a quantum analog of the technique of Bellare and Rogaway [BR93] and demonstrating how to simulate a random oracle without any additional computational assumptions. Specifically, we:

- Describe new ways to argue that quantum algorithms cannot distinguish between two distributions of oracles.
- Apply these techniques to a new type of distribution of oracles, which we call *semi-constant* distributions, showing that they cannot be distinguished from random oracles.

- Use our results on semi-constant distributions to prove that the random oracle IBE scheme of Gentry et al. [GPV08] is secure against quantum adversaries. The basic idea is to plug the challenge c into a small fraction of inputs to the oracle, making the oracle seen by the adversary a semi-constant distribution. The adversary thus behaves as though the oracle is random. If the adversary happens to use any of the inputs in this fraction, we are able to solve c .
- Show that this technique is general by applying it to the random oracle hierarchical IBE schemes of Cash et al. [CHKP10] and Agrawal et al. [ABB10].
- Prove that the generic Full Domain Hash signature scheme [BR93] is secure against quantum adversaries, though this remains a theoretical result until a trapdoor permutation is found that is secure against quantum adversaries.
- Use our techniques and k -wise independent functions to solve the problem of simulating quantum random oracles, thus making the above results unconditional.

1.2 Related Work

Quantum random oracles have been used in several prior works. For example, Bennett et al. [BBBV97] prove several quantum complexity results, including a proof that quantum computers cannot solve all of NP, relative to a random oracle. In a cryptographic setting, quantum random oracles have been used by Aaronson [Aar09] to construct quantum money and by Brassard and Salvail [BS08] and Brassard et al. [BHK⁺11] to construct quantum analogs of Merkle’s Puzzles.

Boneh et al. [BDF⁺11] give a random oracle scheme that is secure against quantum adversaries with classical access to the oracle, but is insecure once the adversary has quantum access to the oracle. Despite this result, they show that there are circumstances in which a classical random oracle security reduction can also be used in the quantum setting. However, their proof techniques do not apply to the schemes analyzed in this paper. Additionally, their proof simulates the random oracle using a pseudorandom function (PRF) that is secure against quantum adversaries, hence making their results conditional on the existence of a so-called quantum-secure PRF. We note that Zhandry [Zha12b] shows that such PRFs can be built from quantum-secure pseudorandom generators, which can, in turn, be built from lattices.

There has also been progress toward converting classical security proofs into quantum proofs outside of the random oracle model. For instance, Unruh [Unr10] shows that classical statistical security in Canetti’s universal composability (UC) model implies quantum statistical security. Hallgren, Smith, and Song [HSS11] give a two-party protocol that is quantum computationally secure in the UC.

2 Preliminaries

A function $\epsilon(n)$ is negligible if it is non-negative and smaller than any inverse polynomial. That is, for any polynomial $p(n)$, $\epsilon(n) < 1/p(n)$ for all sufficiently large n .

A probabilistic polynomial time (PPT) algorithm is a classical randomized algorithm that runs in time polynomial in the size of its input. We also call such algorithms efficient.

2.1 Weight Assignments

A weight assignment on a set \mathcal{X} is a function $D : \mathcal{X} \rightarrow \mathbb{R}$ such that $\sum_{x \in \mathcal{X}} D(x) = 1$. We sometimes write $\Pr_D[\text{event}]$ to represent the sum of the weights of all outcomes consistent with that event. A distribution on X is a weight-assignment D such that $D(x) \geq 0$ for all $x \in \mathcal{X}$. If D is a distribution, way that x occurs with probability $D(x)$. Let $U_{\mathcal{X}}$ denote the uniform distribution over \mathcal{X} . That is, $U_{\mathcal{X}}(x) = 1/|\mathcal{X}|$. When the set \mathcal{X} is clear, we may omit the subscript.

We define the distance between two weight assignments D_1 and D_2 over a set \mathcal{X} as

$$|D_1 - D_2| = \sum_{x \in \mathcal{X}} |D_1(x) - D_2(x)|$$

If $|D_1 - D_2| \leq \epsilon$, we say that D_1 and D_2 are ϵ -close.

Given a set of weight assignments D_y over \mathcal{X} , indexed by $y \in \mathcal{Y}$, and a weight assignment D over \mathcal{Y} , we can define a weight assignment D' over \mathcal{X} where

$$D'(x) = \sum_{y \in \mathcal{Y}} D(y) D_y(x)$$

We write $D' = \sum_{y \in \mathcal{Y}} D(y) D_y$. Say each of the D_y s are actually distributions. Given two weight assignments D_1 and D_2 over \mathcal{Y} , define $D'_1 = \sum_{y \in \mathcal{Y}} D_1(y) D_y$ and $D'_2 = \sum_{y \in \mathcal{Y}} D_2(y) D_y$. Then:

$$\begin{aligned} |D'_1 - D'_2| &= \sum_{x \in \mathcal{X}} \left| \sum_{y \in \mathcal{Y}} (D_1(y) - D_2(y)) D_y(x) \right| \\ &\leq \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} |(D_1(y) - D_2(y)) D_y(x)| \\ &= \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} |D_1(y) - D_2(y)| D_y(x) \\ &= \sum_{y \in \mathcal{Y}} |D_1(y) - D_2(y)| = |D_1 - D_2| \end{aligned}$$

Consider the set of functions $H : \mathcal{X} \rightarrow \mathcal{Y}$ for sets \mathcal{X} and \mathcal{Y} , denoted by $\mathcal{H}_{\mathcal{X}, \mathcal{Y}}$. Consider a weight assignment D on $\mathcal{H}_{\mathcal{X}, \mathcal{Y}}$. Let $\mathcal{W} \subseteq \mathcal{X}$. We define the marginal weight assignment $D_{\mathcal{W}}$ of D on $\mathcal{H}_{\mathcal{W}, \mathcal{Y}}$ where the weight of a function $H_{\mathcal{W}} : \mathcal{W} \rightarrow \mathcal{Y}$ is equal to the sum of the weights of all $H \in \mathcal{H}_{\mathcal{X}, \mathcal{Y}}$ that agree with $H_{\mathcal{W}}$ on \mathcal{W} . In other words,

$$D_{\mathcal{W}}(H_{\mathcal{W}}) = \Pr_D[H(w) = H_{\mathcal{W}}(w) \forall w \in \mathcal{W}]$$

We call two weight assignments D_1 and D_2 on $\mathcal{H}_{\mathcal{X}, \mathcal{Y}}$ k -wise equivalent if for all $\mathcal{W} \subseteq \mathcal{X}$ of size k , the marginal weight assignments $D_{1, \mathcal{W}}$ and $D_{2, \mathcal{W}}$ (of D_1 and D_2) over $\mathcal{H}_{\mathcal{W}, \mathcal{Y}}$ are identical.

2.2 Quantum Computation

A quantum algorithm is an algorithm executed on a quantum computer that produces a classical output. Most of the background in quantum computation needed to understand this paper is for the proof of Theorem 3.1. Since this proof appears in the full version [Zha12a], we present the necessary background there. In the meantime, we recall a couple basic facts about quantum computation, and refer the reader to [NC00] for a more thorough discussion.

Fact 1. *Any classical computation can be implemented on a quantum computer.*

Fact 2. *Any function that has an efficient classical algorithm computing it can be implemented efficiently as a quantum-accessible oracle.*

Fact 3. *Given a quantum algorithm A with oracle access to an oracle O , each oracle O defines a probability distribution of the outputs of A . Hence, any weight assignment of oracles leads to a weight assignment of outputs of A , and if two weight assignments D_1 and D_2 are a distance ϵ apart, the weight assignment of the outputs of A under these distributions are a distance at most ϵ apart.*

2.3 Cryptographic Primitives

Here we briefly outline a few cryptographic primitives, and refer to the full version [Zha12a] for details. All primitives depend on a security parameter n .

An encryption scheme E is a triple of PPT algorithms $(E.Gen, E.Enc, E.Dec)$, where $E.Gen(1^n)$ generates secret/public keys (sk, pk) , $E.Enc_{pk}$ encrypts a message, and $E.Dec_{sk}$ decrypts a ciphertext. We use the indistinguishability under chosen plaintext attack (IND-CPA) notion of security [GM84].

An identity-based encryption (IBE) scheme IBE is a 4-tuple of PPT algorithms $(IBE.Gen, IBE.Extract, IBE.Enc, IBE.Dec)$ where $IBE.Gen(1^n)$ generates master secret/public keys (msk, mpk) , $IBE.Extract_{msk}$ generates secret keys for given identities, $IBE.Enc_{mpk}$ encrypts a message to an identity, and $IBE.Dec$ decrypts a ciphertext sent to an identity by using the corresponding secret key. We use the indistinguishability under chosen plaintext attack (IND-ID-CPA) notion of security [BF01].

A signature scheme $S = (S.Gen, S.Sign, S.Ver)$ is a triple of PPT algorithms where $S.Gen(1^n)$ generates secret/public keys (sk, pk) , $S.Sign_{sk}$ signs a message, and $S.Ver_{pk}$ verifies a signature. We use the existential unforgeability under chosen message attack (UF-CMA) notion of security [GMR88].

A pre-image sampleable function (PSF) is a quadruple of algorithms $F = (F.Gen, F.Sample, f, f^{-1})$ where $F.Gen$ generates private/public keys (sk, pk) , f_{pk} is a function, $F.Sample$ samples x from a distribution D such that $f_{pk}(x)$ is uniform, and $f_{sk}^{-1}(y)$ samples from D conditioned on $f_{pk}(x) = y$. For security, we use the notion of one-wayness.

A trapdoor permutation (TDP) is a special case of a PSF where f is bijective and $F.Sample$ simply returns a random element in the domain of f_{pk} . Since $F.Sample$ is already defined, we omit it from the specification.

2.4 Random Oracle Model

In the Random Oracle Model, we assume the existence of a random function H , and give all parties oracle access to this function. If A makes queries to an oracle H , we denote this as A^H . If a system \mathcal{S} uses a random oracle in its specification, we denote this as \mathcal{S}^H . The algorithms comprising any cryptographic protocol can use H , as can the adversary. Thus we modify the security games for all cryptographic systems to allow the adversary to make random oracle queries.

When a random oracle scheme is implemented, some suitable hash function H is included in the specification. Any algorithm (adversary included) now replaces oracle queries with evaluations of this hash function. In the quantum setting, because a quantum algorithm can evaluate H on an arbitrary superposition, we must allow the quantum adversary to make quantum queries to the random oracle. We call this the quantum random oracle model.

3 Distinguishing Oracles with Quantum Queries

In this section, we give some tools for arguing that quantum algorithms cannot distinguish between two distributions of oracles. The trivial way to bound the ability of any quantum algorithm to distinguish two oracle distributions D_1 and D_2 is to bound the distance between D_1 and D_2 themselves. However, in many cases, the distributions we are interested in are in fact very far apart, so we need some new techniques for arguing indistinguishability. First, we show that k -wise equivalent distributions are indistinguishable, using techniques similar to that of Meyer and Pommersheim [MP11]:

Theorem 3.1. *Let A be a quantum algorithm making q quantum queries to an oracle $H : \mathcal{X} \rightarrow \mathcal{Y}$. If we draw H from some weight assignment D , then for every z , the quantity $\Pr_{H \leftarrow D}[A^H() = z]$ is a linear combination of the quantities $\Pr_{H \leftarrow D}[H(x_i) = r_i \forall i \in \{1, \dots, 2q\}]$ for all possible settings of the x_i and r_i .*

This is proved in the full version [Zha12a], and immediately implies two important facts:

- If two weight assignments on oracles, D_1 and D_2 , are $2q$ -wise equivalent, then any q query quantum algorithm behaves the same under both weight assignments, since for all $2q$ pairs (x_i, r_i) , $\Pr_{H \leftarrow D_1}[H(x_i) = r_i \forall i \in \{1, \dots, 2q\}] = \Pr_{H \leftarrow D_2}[H(x_i) = r_i \forall i \in \{1, \dots, 2q\}]$.
- If the quantities $\Pr_{H \leftarrow D}[H(x_i) = r_i \forall i \in \{1, \dots, 2q\}]$ are low-degree polynomials in some parameter λ , then so is $\Pr_{H \leftarrow D}[A^H() = z]$, since it is a linear combination of the $\Pr_{H \leftarrow D}[H(x_i) = r_i \forall i \in \{1, \dots, 2q\}]$. Moreover, $\Pr_{H \leftarrow D}[A^H() \in S]$ for any set S is also a low-degree polynomial.

We use the second fact to prove the following theorem:

Theorem 3.2. *Fix q , and let D_λ be a family of distributions on $\mathcal{H}_{\mathcal{X}, \mathcal{Y}}$ indexed by $\lambda \in [0, 1]$. Suppose there are integers d and Δ such that for every $2q$ pairs $(x_i, r_i) \in \mathcal{X} \times \mathcal{Y}$, the function $p(\lambda) = \Pr_{H \leftarrow D_\lambda}[H(x_i) = r_i \forall i \in \{1, \dots, 2q\}]$ satisfies:*

- p is a polynomial in λ of degree at most d .
- $p^{(i)}(0)$, the i th derivative of p at 0, is 0 for each $i \in \{1, \dots, \Delta - 1\}$.

Then any quantum algorithm A making q quantum queries can only distinguish D_λ from D_0 with probability at most $\frac{4^\Delta}{(2\Delta)!} \lambda^\Delta d^{2\Delta}$.

Proof. By the assumptions of the theorem, for any $2q$ pairs (x_i, r_i) , the quantity $\Pr_{H \leftarrow D_\lambda}[H(x_i) = r_i \forall i \in \{1, \dots, 2q\}]$ is a polynomial of degree d in λ with the first $\Delta - 1$ derivatives at 0 being 0. By Theorem 3.1, for a q -query quantum algorithm A , $\Pr_{H \leftarrow D_\lambda}[A^H() = z]$ is a linear combination of these values. Thus, for any z , $\Pr_{H \leftarrow D_\lambda}[A^H() = z]$ is also a polynomial in λ of degree d with the first $\Delta - 1$ derivatives at 0 being 0.

Now, suppose that A distinguishes D_λ from D_0 with probability $\epsilon(\lambda)$. That is

$$\sum_z \left| \Pr_{H \leftarrow D_\lambda}[A^H() = z] - \Pr_{H \leftarrow D_0}[A^H() = z] \right| = \epsilon(\lambda) .$$

Let \mathcal{Z}_λ be the set of z such that z is a more likely output under D_λ than D_0 . That is, $\Pr_{H \leftarrow D_\lambda}[A^H() = z] > \Pr_{H \leftarrow D_0}[A^H() = z]$. It is not difficult to show that

$$\Pr_{H \leftarrow D_\lambda}[A^H() \in \mathcal{Z}_\lambda] - \Pr_{H \leftarrow D_0}[A^H() \in \mathcal{Z}_\lambda] = \epsilon(\lambda)/2 .$$

Fix λ_0 , and consider the quantity

$$p_{\lambda_0}(\lambda) \equiv \Pr_{H \leftarrow D_\lambda}[A^H() \in \mathcal{Z}_{\lambda_0}] = \sum_{z \in \mathcal{Z}_{\lambda_0}} \Pr_{H \leftarrow D_\lambda}[A^H() = z] .$$

Then $p_\lambda(\lambda) - p_\lambda(0) = \epsilon(\lambda)/2$. Further, for each λ_0 , p_{λ_0} is a degree- d polynomial in λ such that $p_{\lambda_0}^{(i)}(0) = 0$ for $i \in \{1, \dots, \Delta - 1\}$. It also lies in the range $[0, 1]$ for all $\lambda \in [0, 1]$, so we can use an inequality by the Markov brothers [DS41] to bound the Δ -th derivative for all $\lambda \in [0, 1]$:

$$\left| p_{\lambda_0}^{(\Delta)}(\lambda) \right| \leq \frac{4^\Delta \Delta!}{2(2\Delta)!} d^{2\Delta}$$

Thus we can bound $p_{\lambda_0}(\lambda) \leq p_{\lambda_0}(0) + \frac{4^\Delta}{2(2\Delta)!} \lambda^\Delta d^{2\Delta}$. Setting $\lambda_0 = \lambda$, we get

$$\epsilon(\lambda) = 2(p_\lambda(\lambda) - p_\lambda(0)) \leq \frac{4^\Delta}{(2\Delta)!} \lambda^\Delta d^{2\Delta} .$$

□

We conclude this section with another general approach toward arguing that no quantum algorithm can distinguish between two oracle distributions D_1 and D_2 : construct a weight assignment $D_{1,2}$ which is $2q$ -wise equivalent to D_1 (so that the behaviors of A under D_1 and $D_{1,2}$ are the same) and that is a distance at most ϵ from D_2 (so that the behaviors of A under $D_{1,2}$ and D_2 are ϵ -close). Then the behaviors of A under D_1 and D_2 are also ϵ -close. We formalize this latter idea with a new pseudometric on oracle distributions:

Definition 3.3. Let D_1 and D_2 be two weight assignments on oracles. Fix k . Then define the pseudometric $|D_1 - D_2|_{(k)}$ as the minimum of $|D_{1,2} - D_2|$ over all weight assignments $D_{1,2}$ that are k -wise equivalent D_1 .

Theorem 3.4. The minimum $D_{1,2}$ in Definition 3.3 exists, and $|D_1 - D_2|_{(k)}$ defines a pseudometric on weight assignments on oracles.

Proof. If D is k -wise equivalent to D_1 and a distance ϵ from D_2 , we say that D is a witness to $|D_1 - D_2|_{(k)} \leq \epsilon$. If $\epsilon = |D_1 - D_2|_{(k)}$, we call D an optimal witness.

First, we argue that an optimal witness exists: We see that D_1 itself is a witness to $|D_1 - D_2|_{(k)} \leq |D_1 - D_2|$. Thus to find $D_{1,2}$, we are minimizing $|D_{1,2} - D_2|$ (a continuous function in $D_{1,2}$) over the set of assignments $D_{1,2}$ that are k -wise similar to D_1 (a closed set) and at most $|D_1 - D_2|$ away from D_2 (a compact set). Thus we are minimizing over a compact set, so the minimum is actually obtained.

We now show that $|D_1 - D_2|_{(k)}$ satisfies the properties of a pseudometric:

- $|D_1 - D_1|_{(k)} = 0$, since D_1 is a witness to $|D_1 - D_1|_{(k)} \leq 0$, and $|D_1 - D_2|_{(k)}$ is non-negative for all weight assignments D_1 and D_2 .
- $|D_1 - D_2|_{(k)} = |D_2 - D_1|_{(k)}$. Let $D_{1,2}$ be an optimal witness for $|D_1 - D_2|_{(k)}$. Then define $D_{2,1} = D_1 - D_{1,2} + D_2$. The sums of all weights add to 1, so this is a weight assignment. Further, it is easy to see that $|D_{2,1} - D_1| = |D_2 - D_{1,2}| = |D_1 - D_2|_{(k)}$. Further, since D_1 is k -wise equivalent to $D_{1,2}$, when looking at the marginal weight assignment over any k inputs, D_1 and $D_{1,2}$ cancel, leaving $D_{2,1}$ to be k -wise equivalent to D_2 . Thus, $D_{2,1}$ is a witness to $|D_2 - D_1|_{(k)} \leq |D_1 - D_2|_{(k)}$. We can perform the same argument in reverse to get the other inequality.
- $|D_1 - D_3|_{(k)} \leq |D_1 - D_2|_{(k)} + |D_2 - D_3|_{(k)}$. Let $D_{1,2}$ and $D_{2,3}$ be the optimal witnesses for $|D_1 - D_2|_{(k)}$ and $|D_2 - D_3|_{(k)}$. Define $D_{1,3} = D_{1,2} - D_2 + D_{2,3}$. It is easy to see that D_3 is a weight assignment and that $|D_{1,3} - D_3| \leq |D_{1,2} - D_2| + |D_{2,3} - D_3|$. It is also not difficult to show that, since $D_{2,3}$ is k -wise equivalent to D_2 , $D_{1,3}$ must be k -wise equivalent to $D_{1,2}$, which is k -wise equivalent to D_1 . Hence $D_{1,3}$ is a witness for the identity in question. \square

Observe that $|D_1 - D_2|_{(k)} = 0$ if and only if D_1 is k -wise equivalent to D_2 . Thus, we can rephrase Theorem 3.1 as follows: if $|D_1 - D_2|_{(2q)} = 0$, then no quantum algorithm making q quantum queries can distinguish the oracle weight assignments D_1 and D_2 . We now generalize this to $|D_1 - D_2|_{(2q)} \neq 0$:

Corollary 3.5. Let D_1 and D_2 be two oracle distributions. Then the output weight assignment of any q -query quantum algorithm A under D_1 and D_2 are $|D_1 - D_2|_{(2q)}$ -close.

Proof. Let $D_{1,2}$ be an optimal witness for $|D_1 - D_2|_{(2q)}$. $D_{1,2}$ is $2q$ -wise equivalent to D_1 , so by Theorem 3.1, the behaviors of the A under D_1 and $D_{1,2}$ are identical. The behavior of A under D_2 is $|D_{1,2} - D_2| = |D_1 - D_2|_{(2q)}$ -close to that under $D_{1,2}$, and hence D_1 . \square

We note that a slightly weaker version of Theorem 3.2 can also be proved using the above ideas. Roughly, we pick some constants λ_i and find functions $a_i(\lambda)$ such that

$$p(\lambda) = \left(1 - \sum_{i=0}^{d-\Delta} a_i(\lambda)\right) p(0) + \sum_{i=0}^{d-\Delta} a_i(\lambda) p(\lambda_i)$$

for any polynomial p of degree d such that $p^{(i)}(0) = 0$ for $i \in \{1, \dots, \Delta - 1\}$. The $a_i(\lambda)$ are basically modified Lagrange interpolating polynomials. Next, we define a new weight assignment

$$D'_\lambda = \left(1 - \sum_{i=0}^{d-\Delta} a_i(\lambda)\right) D_0 + \sum_{i=0}^{d-\Delta} a_i(\lambda) D_{\lambda_i} .$$

By defining D'_λ in this way, D'_λ turns out to be $2q$ -wise equivalent to D_λ . Because D_{λ_i} are all distributions, the distance between D'_λ and D_0 is at most $2 \sum_{i=0}^{d-\Delta} |a_i(\lambda)|$, thus showing that $|D_\lambda - D_0|_{(2q)}$ also satisfies this bound. By picking the right values for λ_i , we can get this bound to be small. We can then use Corollary 3.5 to bound the distance between the output distributions of the q -query quantum algorithm A .

4 Semi-constant Distributions

In this section, we define a class of distributions on oracles in $\mathcal{H}_{\mathcal{X}, \mathcal{Y}}$ called semi-constant distributions. Our motivation for these distributions is to support quantum Full Domain Hash-type arguments, where a random value is inserted into a small but significant fraction of oracle inputs. The following definition captures the essence of the this idea:

Definition 4.1. Define SC_λ , the semi-constant distribution, as the distribution over $\mathcal{H}_{\mathcal{X}, \mathcal{Y}}$ resulting from the following process:

- First, pick a random element y from \mathcal{Y} .
- For each $x \in \mathcal{X}$, do one of the following:
 - With probability λ , set $H(x) = y$. We call x a distinguished input to H .
 - Otherwise, set $H(x)$ to be a random element in \mathcal{Y} .

First, notice that SC_0 is the uniform distribution, since in this case, the probability that an input is distinguished is 0. We bound the ability of a quantum algorithm to distinguish between SC_λ and $\text{SC}_0 = U$ using the ideas of Section 3.

Lemma 4.2. Fix k . For each k pairs (x_i, r_i) , the probability $\Pr_{H \leftarrow \text{SC}_\lambda}[H(x_i) = r_i \forall i \in \{1, \dots, k\}]$ is a degree- k polynomial in λ whose first derivative is 0 at $\lambda = 0$.

This, proved in the full version [Zha12a], shows that for any q_H , we can set $d = k = 2q_H$ and $\Delta = 2$ in the assumptions of Theorem 3.2. We immediately get:

Corollary 4.3. *The distribution of outputs of a quantum algorithm making q_H queries to an oracle drawn from SC_λ is at most a distance $\frac{8}{3}q_H^4\lambda^2$ away from the case when the oracle is drawn from the uniform distribution.*

We note that using standard quantum query results, it is possible to prove this corollary for a bound that is first-order in λ . However, we will see in Section 5 that we need a bound that is second-order in λ for our security results.

We do not know if this bound is tight. In the full version [Zha12a], we adapt the collision search algorithm of Brassard et al. [BHT97] to SC_λ , obtaining a distinguishing probability of $\Omega(q_H^3\lambda^2)$.

5 Quantum Security Arguments

In this section, we explore the random oracle proof technique of Bellare and Rogaway [BR93]. In this technique, we plug a challenge c into a randomly chosen hash query, and hope that the adversary uses c to break the system. If so, we can use the adversary to solve the underlying problem.

In the quantum setting, the straight-forward application of this technique breaks down. The adversary can now query the hash function on a superposition of exponentially many inputs. If we plug c to one of those inputs, the probability that the adversary uses c is exponentially small.

We instead plug c into some small but significant fraction of the inputs, so that the oracle is distributed according to SC_λ . As we have shown, a quantum algorithm cannot detect that this oracle is not random, but now the adversary uses c with significant probability.

5.1 Identity-Based Encryption from Lattices

Here we prove the security of the IBE scheme from Gentry et al. [GPV08]. Their scheme is constructed from an encryption scheme $\mathbf{E} = (\mathbf{E}.\text{Gen}, \mathbf{E}.\text{Enc}, \mathbf{E}.\text{Dec})$, for which there exists a trapdoor allowing the computation of secret keys from public keys.

More specifically, let $\mathbf{F} = (\mathbf{F}.\text{Gen}, \mathbf{F}.\text{Sample}, f, f^{-1})$ be a pre-image sampleable function (PSF). Suppose $\mathbf{E}.\text{Gen}(1^n)$ works as follows: generate $(\text{msk}, \text{mpk}) \leftarrow \mathbf{F}.\text{Gen}(1^n)$. Then, sample $\text{sk} \leftarrow \mathbf{F}.\text{Sample}(1^n)$, and compute $\text{pk} = f_{\text{mpk}}(\text{sk})$. Output $(\text{sk}, (\text{pk}, \text{mpk}))$.

Gentry et al. give such an encryption scheme based on the hardness of lattice problems. Their security reduction treats the adversary as a black box, so their proof holds in the quantum setting. They then prove the security of the IBE scheme $\text{IBE} = (\text{IBE}.\text{Gen} = \mathbf{F}.\text{Gen}, \text{IBE}.\text{Extract}^H, \text{IBE}.\text{Enc}^H, \text{IBE}.\text{Dec})$ where H maps identities to public keys of \mathbf{E} and:

- $\text{IBE}.\text{Extract}_{\text{msk}}(\text{id})^H = f_{\text{msk}}^{-1}(H(\text{id}))$
- $\text{IBE}.\text{Enc}_{\text{mpk}}(\text{id}, m)^H = \mathbf{E}.\text{Enc}_{H(\text{id})}(m)$
- $\text{IBE}.\text{Dec}_{\text{sk}_{\text{id}}}(c) = \mathbf{E}.\text{Dec}_{\text{sk}_{\text{id}}}(c)$

Theorem 5.1. *Let \mathbf{E} and \mathbf{F} be as above, and suppose that \mathbf{E} is quantum IND-CPA-secure. If we model H as a random oracle, then IBE is quantum IND-ID-CPA-secure.*

Proof. Let A be a quantum adversary making q_H hash queries, q_E extract queries, that breaks IBE with advantage ϵ .

Let Game_0 be the standard attack game for IBE: the challenger generates (msk, mpk) from IBE.Gen , and sends mpk to the adversary. The adversary can make (classical) extraction queries on identities id_i , to which the challenger responds with $\text{IBE.Extract}_{\text{msk}}^H(\text{id}_i)$, and (quantum) hash queries to the random oracle H . A then produces an identity id^* , along with messages m_0 and m_1 . The challenger chooses a random bit b , and responds with $\text{IBE.Enc}_{\text{mpk}}(\text{id}, m_b)$. A is allowed to make more extraction and hash queries, and produces a bit b' . A wins if $b' = b$ and for all i , $\text{id}_i \neq \text{id}^*$. By definition, this happens with probability $\frac{1}{2} + \epsilon$.

Assume without loss of generality that A never asks for the secret key corresponding to id^* . Let A' be the following algorithm that makes quantum queries to an oracle H , and simulates the interaction between A and the challenger: generate (msk, mpk) from IBE.Gen , send mpk to A . Run A , forwarding all hash queries to H , and answering extraction queries using IBE.Extract^H . Similarly, answer the challenge query by choosing a random b and encrypting m_b to the identity id generated by A . The output of A' is (c, id^*, Q) where $c = b \oplus b'$ and Q is the list of extraction queries made by A . We can now think of Game_0 as follows: run A' with a random oracle H to obtain (c, id^*, Q) . Report that the game is won if and only if $c = 0$. The number of queries to H made by A' is q_H for direct queries, q_E for queries through the extract algorithm, and 1 for the encryption of m_b , for a total of $q_H + q_E + 1$ queries.

Let $\lambda \in (0, 1)$ to be chosen later. Let \mathcal{M} be the identity space. Define Game_1 as follows: first, pick a subset \mathcal{X} of \mathcal{M} where each $\text{id} \in \mathcal{M}$ is put in \mathcal{X} with independent probability λ . Run A' as before to obtain (c, id^*, Q) . However, now we decide if the game is won as follows: if $\text{id}^* \in \mathcal{X}$ and $Q \cap \mathcal{X} = \emptyset$, then the game is won if $c = 0$. Otherwise, flip a random coin to decide if the game is won. By our assumption that A never asks for a secret key for the challenge identity, we must have that $\text{id}^* \notin Q$. Therefore, every identity in Q , as well as id^* , has an independent probability λ of being in \mathcal{X} . Thus, the conditions to not flip a coin are met with probability at least $\lambda(1 - \lambda q_E)$, independent of the probability that $c = 0$. Thus, Game_1 is won with probability at least $\frac{1}{2} + \lambda(1 - \lambda q_E)\epsilon \geq \frac{1}{2} + \lambda\epsilon - q_E\lambda^2$.

Let Game_2 be Game_1 , except that we choose a random pk in the public key space of \mathbf{E} , and set $H(x) = \text{pk}$ for all x in \mathcal{X} , and choose $H(x)$ randomly for all other x . H is now distributed according to SC_λ .

Claim. There exists a polynomial $\ell(\cdot, \cdot)$ such that A' wins Game_2 with probability at least

$$\frac{1}{2} + \lambda\epsilon - \frac{\ell(q_H, q_E)}{4}\lambda^2.$$

We defer the precise proof to the full version [Zha12a], and instead give an informal justification: by Corollary 4.3, the output distribution of A' in Game_2 is at most a distance $\frac{8}{3}(q_H + q_E + 1)^4 \lambda^2$ from that of Game_1 . Thus, we would expect A' to win Game_2 with probability at least $\frac{1}{2} + \lambda\epsilon - (q_E + \frac{8}{3}(q_H + q_E + 1)^4) \lambda^2$, which has the desired form. Of course, deciding if A' wins Game_2 depends on the set of distinguished inputs \mathcal{X} , so a more careful analysis is required to justify the claim.

Let us now restate Game_2 in terms of A : pick \mathcal{X} and H as above, and perform the standard attack game for IBE using H as the random oracle. When A produces b' , first check that no extraction queries were in \mathcal{X} and that the identity id^* produced by A is in \mathcal{X} . If so, A wins if and only if $b = b'$. Otherwise, if the check fails, flip a coin to decide if A wins. We can now make the following observation: once an extraction query in \mathcal{X} is made or the challenge query with $\text{id}^* \notin \mathcal{X}$ is made, we can abort running A , since we no longer need A to determine if the game is won.

Now we describe an adversary B that breaks E. Assume B has quantum access to two random oracles O_1 and O_2 . O_1 maps identities to randomness used by Sample . O_2 maps identities to bits, outputting 1 with probability λ . In Section 6, we will show how to efficiently simulate these oracles. On input (mpk, pk) , B works as follows:

- Send mpk to A and simulate A , playing the role of challenger to A .
- Construct a (quantum) oracle H such that

$$H(\text{id}) = \begin{cases} \text{pk} & \text{if } O_2(\text{id}) = 1 \\ f_{\text{mpk}}(\text{Sample}(\cdot; O_1(\text{id}))) & \text{otherwise} \end{cases}$$

where $\text{Sample}(\cdot; r)$ means run Sample with randomness r .

- When A asks for the secret key for id_i , compute $O_2(\text{id}_i)$. If the result is 1, output a random bit and abort. Otherwise, respond with $\text{sk}_i = \text{Sample}(\cdot; O_1(\text{id}_i))$.
- When A produces the challenge query (id^*, m_0, m_1) , check if $O_2(\text{id}^*) = 1$. If so, send (m_0, m_1) to B 's challenger. Otherwise, output a random bit and abort.
- When the challenger responds with a ciphertext, send it to A .
- When A outputs a bit b' , output the same bit.

Let \mathcal{X} be the set of identities id such that $O_2(\text{id}) = 1$. We can then see that the abort conditions are equivalent to Game_2 . For each extract query id_i that succeeds, we have that $O_2(\text{id}_i) = 0$, so

$$f_{\text{mpk}}(\text{sk}_i) = f_{\text{mpk}}(\text{Sample}(\cdot; O_1(\text{id}_i))) = H(\text{id}_i)$$

Thus sk_i is a correct secret key for id_i , and it is distributed correctly (since it is a random pre-image of $H(\text{id}_i)$). If B does not abort on the challenge query, it means $O_2(\text{id}^*) = 1$, so $H(\text{id}^*) = \text{pk}$. Also, we can see that H is distributed according to SC_λ . Therefore, if B aborts, we win with probability $\frac{1}{2}$, and if we do not abort, we win if and only if A wins. Thus, the view of A as a subroutine

of B is the same as in Game_2 , and B wins with the same probability that A does. Therefore, B wins with probability at least $\frac{1}{2} + \lambda\epsilon - \frac{\ell(q_H, q_E)}{4}\lambda^2$.

The advantage of B is at least $\lambda\epsilon - \ell(q_H, q_E)\lambda^2/4$. We can now choose λ : the maximum advantage occurs when $\lambda = 2\epsilon/\ell(q_H, q_E)$, which gives

$$\text{Adv}_B \geq \frac{\epsilon^2}{\ell(q_H, q_E)}$$

Thus, if A had non-negligible advantage, so does B . \square

5.2 Hierarchical IBE from Lattices

In this section, we show the general idea behind adapting the techniques above to proving the security of the hierarchical identity-based encryption (HIBE) schemes of Agrawal et al. [ABB10] and Cash et al. [CHKP10].

In a HIBE scheme, identities are structured as a tree, with the identity of any node containing the identity of its parent as a proper prefix. Any node on the tree can produce private keys for any nodes in the subtree rooted at that node. We allow an adversary to adaptively take control of an arbitrary number of nodes in the tree (and thus the subtrees rooted there). An HIBE scheme is secure if the adversary cannot decrypt messages encrypted to an identity id^* of the adversary's choice but not under its control.

In [ABB10], the random oracle scheme has an oracle H that maps identities to some random quantities. The reduction has the following high-level structure:

- Guess which level w of the tree contains the identity id^*
- For each level i , generate some random quantities R_i .
- For each level i , simulate a separate random oracle for identities at that level. For i , guess which query number q_i will contain the hash of the level- i parent of id^* .
- Answer the j th random oracle query at level i on $\text{id}_{i,j}$ as follows: if $j = q_i$, output R_i . Otherwise, output a random value.
- Answer secret key queries on id in some special way, but fail if for all prefixes id_i of id , $\text{id}_i = \text{id}_{i,q_i}$. That is, fail if $H(\text{id}_i) = R_i$.
- When the adversary generates the identity id^* , we succeed if both the adversary succeeds and if id^* is at level w and all prefixes id_i^* of id^* satisfy $\text{id}_i^* = \text{id}_{i,q_i}$.

We now show how to prove security by repeatedly applying the arguments of Theorem 5.1. Basically, we iterate over levels i , and insert R_i into a λ_i fraction of identities at that level. In iteration i , we say the adversary wins if it won in the previous iteration, the level- i prefix of the chosen identity id^* is in the λ_i fraction of distinguished identities (that is, $H(\text{id}_i^*) = R_i$), and no signature query is. Let $\ell = \ell(q_H, q_E)$, where $\ell(\cdot, \cdot)$ is the polynomial from Theorem 5.1. If the iteration i advantage is ϵ_i , then using the same techniques as in Theorem 5.1, we can set λ_i so that

$$\epsilon_i \geq \frac{\epsilon_{i-1}^2}{\ell}$$

We will say that in iteration 0, the adversary wins if it normally would win and we guessed which level id^* belonged to correctly. That is, $\epsilon_0 = \epsilon/d$, where ϵ is the adversary's advantage in the standard game. This gives us a total advantage after iteration d of at least

$$\frac{(\epsilon/d)^{2^d}}{\ell^{(2^d-1)}} = \ell \left(\frac{\epsilon}{d\ell} \right)^{2^d}$$

Notice that the dependence on d is doubly-exponential, whereas in the classical proof it was singly exponential. Thus, for the same security parameters, this proof only works for much smaller depth than the classical proof.

These techniques apply as well to the random oracle HIBE of Cash et al. [CHKP10], though their reduction is a bit more complicated, as there is a second random oracle G which needs to be handled in a similar way.

5.3 Signatures from Trapdoor Permutations

Here we discuss the security of the Full Domain Hash (FDH) signature scheme:

Definition 5.2 (FDH Signatures). Let $F = (F.\text{Gen}, f, f^{-1})$ be a trapdoor permutation, and a hash function H that maps messages to images of f , let $S^H = (S.\text{Gen} = F.\text{Gen}, S.\text{Sign}^H, S.\text{Ver}^H)$ where:

- $S.\text{Sign}_{\text{sk}}^H(m) = f_{\text{sk}}^{-1}(H(m))$
- $S.\text{Ver}_{\text{pk}}^H(m, \sigma) = \begin{cases} \text{accept} & \text{if } f_{\text{pk}}(\sigma) = H(m) \\ \text{reject} & \text{otherwise} \end{cases}$

We now state the main theorem of this section:

Theorem 5.3. Let F be a quantum one-way trapdoor permutation. If we model H as a random oracle, then S is quantum UF-CMA-secure.

The proof of this theorem is very similar to that of Theorem 5.1, and appears in the full version [Zha12a].

6 Simulating Random Oracles

In this last section, we explain how to efficiently simulate random oracles. All quantum random oracle proofs to date require the reduction algorithm to have quantum access to a random oracle, but a truly random oracle requires exponentially many bits of randomness to construct. We show that this is not a problem:

Theorem 6.1. Any quantum algorithm A making quantum queries to random oracles O_i can be efficiently simulated by a quantum algorithm B , which has the same output distribution, but makes no queries.

Proof. We construct an algorithm B which simulates A , and answers queries to oracle O_i with evaluations of efficient functions f_i . Boneh et al. [BDF⁺11] use pseudorandom functions (PRF) for the f_i . At first glance, this seems like the only option, as we need a function f_i that cannot be distinguished from random.

Notice, however, that f_i need not be secure against all adversaries, just the adversary we are simulating. We know that our adversary makes q_i queries to oracle O_i , so it suffices to have f_i be PRFs secure for up to q_i quantum queries. In the classical setting, q_i -wise independent functions (functions that are q_i -wise equivalent to a random function) serve as *perfectly* secure PRFs for up to q_i classical queries. We could hope that something similar holds in the quantum world: indeed, according to Theorem 3.1, if f_i is $2q_i$ -wise equivalent to a random function, then the behavior of our adversary is the same when the oracle is random and when it is f_i . Thus if the f_i are $2q_i$ -wise independent, algorithm A , as a subroutine of B , behaves identically to the case where A is given truly random oracles. Hence, the output distribution of B is identical to that of A .

Efficient constructions of k -wise independent functions have been known for some time [Jof74, KM94], and they have been used extensively in the derandomization literature [Lub85, ABI86, KM93]. One common approach to construct a k -wise independent function f from \mathcal{X} to \mathcal{Y} is to assume that $N = |\mathcal{Y}|$ is a prime power and interpret \mathcal{Y} as the field \mathbb{F}_N . Then define a matrix C with the following properties:

- The entries are elements in \mathbb{F}_N .
- There are $|\mathcal{X}|$ rows and some small number r of columns.
- Each subset of k rows is linearly independent.

One such example is the Vandermonde matrix, which is used by Alon et al. [ABI86]. To define the function f , we then pick a random vector v in \mathbb{F}_N^r . $f(x)$ is then the x th element of the vector Cv . Since any k rows of C are linearly independent, any k elements of Cv are independent, and hence f is k -wise independent. The key to making this efficient is that to compute $f(x)$, we only need the x th row of C , which we can compute on the fly. \square

Hence, we can simulate O_1 from Theorem 5.1. To simulate O_2 , which outputs a bit such that $O_2(x) = 1$ with probability λ , approximate λ by some rational number a/b where b is a prime power, and construct a k -wise independent function f' with range $\mathcal{Y} = \{1, \dots, b\}$. Then set

$$f(x) = \begin{cases} 1 & \text{if } f'(x) \leq a \\ 0 & \text{otherwise} \end{cases}$$

7 Conclusion

We have shown how to adapt certain classical random oracle arguments to the quantum random oracle model. Specifically, we gave quantum security proofs for the IBE scheme of Gentry et al. [GPV08] and the Full Domain Hash signature scheme. We achieved this by defining a distribution of oracles, called

semi-constant distributions, and showing that such oracles cannot be distinguished from a random oracle by a quantum adversary. We also show how these techniques can be applied to the random oracle HIBE schemes of Cash et al. [CHKP10] and Agrawal et al. [ABB10]. Lastly, we have shown how to remove the need for quantum-secure pseudorandom functions from prior work.

Although we have made progress toward converting classical random oracle proofs into quantum proofs, there is still much work to be done. First, reductions using our technique result in an algorithm whose advantage is the square of the advantage of the underlying adversary, divided by some other factors. While improving the reduction to first-order in the adversary's advantage would still be of interest for the IBE scheme and Full Domain Hash, it would not affect the qualitative security statements. However, for the HIBE schemes, this second-order behavior results in an exponential weakening of the security relative to the classical reduction for deep identity trees. Therefore, to get the same qualitative security statements as in the classical world, our technique needs to be refined to make the reduction first-order in ϵ .

Further work lies in proving the security of other Random Oracle Model schemes. For example, the construction of signatures from identification protocols by Fiat and Shamir [FS87], though similar to the proofs in this paper, still needs a quantum proof. The difficulty in the Fiat-Shamir reduction is that the plugging step initiates the underlying identification protocol, and there is no obvious quantum analog for this strategy. Also, different types of security arguments, such as Fujisaki and Okamoto's generic conversion of weakly secure encryption schemes into a CCA-secure encryption scheme [FO99], still lack a quantum proof of security.

Acknowledgments. We would like to thank Dan Boneh for his guidance and many insightful discussions. This work was supported by NSF and DARPA. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Department of Defense or the U.S. Government.

References

- [Aar09] Aaronson, S.: Quantum Copy-Protection and Quantum Money. In: Proceedings of the 24th Annual IEEE Conference on Computational Complexity, CCC (2009)
- [ABB10] Agrawal, S., Boneh, D., Boyen, X.: Lattice Basis Delegation in Fixed Dimension and Shorter-Ciphertext Hierarchical IBE. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 98–115. Springer, Heidelberg (2010)
- [ABI86] Alon, N., Babai, L., Itai, A.: A Fast and Simple Randomized Parallel Algorithm for the Maximal Independent Set Problem. *Journal of Algorithms* 7(4), 567–583 (1986)
- [BBBV97] Bennett, C.H., Bernstein, E., Brassard, G., Vazirani, U.: Strengths and Weaknesses of Quantum Computing. *SIAM Journal on Computing* 26, 1510–1523 (1997)

- [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)
- [BF01] Boneh, D., Franklin, M.: Identity-Based Encryption from the Weil Pairing. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 213–229. Springer, Heidelberg (2001)
- [BF11] Boneh, D., Freeman, D.M.: Homomorphic Signatures for Polynomial Functions. In: Paterson, K.G. (ed.) EUROCRYPT 2011. LNCS, vol. 6632, pp. 149–168. 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)
- [BHT97] Brassard, G., Høyer, P., Tapp, A.: Quantum Algorithm for the Collision Problem. ACM SIGACT News (Cryptology Column) 28, 14–19 (1997)
- [BR93] Bellare, M., Rogaway, P.: Random Oracles are Practical: A Paradigm for Designing Efficient Protocols. In: Proceedings of the 1st ACM Conference on Computer and Communications Security (CCS), pp. 62–73. ACM (November 1993)
- [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)
- [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)
- [DS41] Duffin, R.J., Schaffer, A.C.: A Refinement of an Inequality of the Brothers Markoff. Trans. Amer. Math. Soc. 44(3), 289–297 (1941)
- [FO99] Fujisaki, E., Okamoto, T.: Secure Integration of Asymmetric and Symmetric Encryption Schemes. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 537–554. Springer, Heidelberg (1999)
- [FS87] Fiat, A., Shamir, A.: How to Prove Yourself: Practical Solutions to Identification and Signature Problems. In: Odlyzko, A.M. (ed.) CRYPTO 1986. LNCS, vol. 263, pp. 186–194. Springer, Heidelberg (1987)
- [GKV10] Dov Gordon, S., Katz, J., Vaikuntanathan, V.: A Group Signature Scheme from Lattice Assumptions. In: Abe, M. (ed.) ASIACRYPT 2010. LNCS, vol. 6477, pp. 395–412. Springer, Heidelberg (2010)
- [GM84] Goldwasser, S., Micali, S.: Probabilistic Encryption. Journal of Computer and System Sciences, 270–299 (1984)
- [GMR88] Goldwasser, S., Micali, S., Rivest, R.L.: A Digital Signature Scheme Secure Against Adaptive Chosen-Message Attacks. SIAM Journal on Computing 17(2), 281–308 (1988)
- [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)
- [Jof74] Joffe, A.: On a Set of Almost Deterministic k -Independent Random Variables. The Annals of Probability 2(1), 161–162 (1974)

- [KM93] Koller, D., Megiddo, N.: Constructing Small Sample Spaces Satisfying Given Constraints. In: Proceedings of the 25th Annual ACM Symposium on Theory of Computing (STOC), pp. 268–277. ACM (1993)
- [KM94] Karloff, H., Mansour, Y.: On Construction of k -Wise Independent Random Variables. In: Proceedings of the 26th Annual ACM Symposium on Theory of Computing (STOC), vol. 17, pp. 564–573 (1994)
- [Lub85] Luby, M.: A Simple Parallel Algorithm for the Maximal Independent Set Problem. In: Proceedings of the 17th Annual ACM Symposium on Theory of Computing (STOC), pp. 1–10. ACM (1985)
- [MP11] Meyer, D., Pommersheim, J.: On the Uselessness of Quantum Queries. Theoretical Computer Science, 1–12 (March 2011)
- [NC00] Nielsen, M.A., Chuang, I.: Quantum Computation and Quantum Information. American Journal of Physics 70(5), 558 (2000)
- [Sho97] Shor, P.W.: Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. SIAM Journal on Computing 26(5), 1484–1509 (1997)
- [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 (February 2012); Full version available at the Cryptology ePrint Archives, <http://eprint.iacr.org/2012/076/>
- [Zha12b] Zhandry, M.: How to Construct Quantum Random Functions (April 2012); Full version available at the Cryptology ePrint Archives, <http://eprint.iacr.org/2012/182/>